

The Upgraded DØ Detector

DØ Collaboration

V.M. Abazov,^{ag} B. Abbott,^{br} M. Abolins,^{bi} B.S. Acharya,^{aa} M. Adams,^{av} T. Adams,^{at}
M. Agelou,^q J.-L. Agram,^r S.H. Ahn,^{ac} M. Ahsan,^{bc} G.D. Alexeev,^{ag} G. Alkhazov,^{ak}
A. Alton,^{bh} G. Alverson,^{bg} G.A. Alves,^b M. Anastasoae,^{af} J. Anderson,^{au} S. Anderson,^{ap}
B. Andrieu,^p R. Angstadt,^{au} V. Anosov,^{ag} Y. Arnoud,^m A. Askew,^{bv} B. Åsman,^{al}
O. Atramentov,^{ba} C. Autermann,^t C. Avila,^g L. Babukhadia,^{bo} F. Badaud,^l A. Baden,^{be}
L. Bagby,^{au} B. Baldin,^{au} P.W. Balm,^{ae} S. Banerjee,^{aa} E. Barberis,^{bg} O. Bardou,^{bg}
W. Barg,^d P. Bargassa,^{bv} P. Baringer,^{bb} C. Barnes,^{an} J. Barreto,^b J.F. Bartlett,^{au}
U. Bassler,^p M.A. Baturitsky,^{ag} D. Bauer,^{ay} B. Baumbaugh,^{az} A. Bean,^{bb} S. Beauceron,^p
M. Begel,^{bn} A. Bellavance,^{bk} S.B. Beri,^z G. Bernardi,^p R. Bernhard,^{47,*} I. Bertram,^{am}
M. Besançon,^q R. Beuselinck,^{an} D. Beutel,^{aw} V.A. Bezzubov,^{aj} P.C. Bhat,^{au} V. Bhatnagar,^z
M. Binder,^x A. Bishoff,^{as} K.M. Black,^{bf} I. Blackler,^{an} G. Blazey,^{aw} F. Blekman,^{ae}
S. Blessing,^{at} D. Bloch,^r U. Blumenschein,^v V. Bodyagin,^{ai} A. Boehnlein,^{au} O. Boeriu,^{az}
T.A. Bolton,^{bc} D. Bonifas,^{au} F. Borchering,^{au} G. Borissov,^{am} K. Bos,^{ae} T. Bose,^{bm}
C. Boswell,^{as} M. Bowden,^{au} A. Brandt,^{bt} R. Brock,^{bi} G. Brooijmans,^{bm} A. Bross,^{au}
N.J. Buchanan,^{at} D. Buchholz,^{ax} M. Buehler,^{av} V. Buescher,^v S. Burdin,^{au} T.H. Burnett,^{bx}
E. Busato,^p D. Butler,^{au} J.M. Butler,^{bf} J. Bystricky,^q L. Canal,^{au} W. Carvalho,^c
B.C.K. Casey,^{bs} N.M. Cason,^{az} H. Castilla-Valdez,^{ad} S. Chakrabarti,^{aa} D. Chakraborty,^{aw}
K.M. Chan,^{bn} A. Chandra,^{aa} D. Chapin,^{bs} F. Charles,^r E. Cheu,^{ap} L. Chevalier,^q
E. Chi,^{au} R. Chiche,^o D.K. Cho,^{bn} R. Choate,^{au} S. Choi,^{as} J.H. Christenson,^{au}
T. Christiansen,^x L. Christofek,^{bb} I. Churin,^{ag} G. Cisco,^{au} D. Claes,^{bk} B. Clément,^r
C. Clément,^{al} Y. Coadou,^e M. Cooke,^{bv} W.E. Cooper,^{au} D. Coppage,^{bb} M. Corcoran,^{bv}
J. Coss,^s A. Cothenet,ⁿ M.-C. Cousinou,ⁿ S. Crépe-Renaudin,^m M. Cristetiu,^{as}
M.A.C. Cummings,^{aw} D. Cutts,^{bs} H. da Motta,^b B. Davies,^{am} G. Davies,^{an} G.A. Davis,^{ax}
W. Davis,^{au} K. De,^{bt} P. de Jong,^{ae} S.J. de Jong,^{af} E. De La Cruz-Burelo,^{ad}
C. De La Taille,^o C. De Oliveira Martins,^c S. Dean,^{ao} F. Déliot,^q P.A. Delsart,^s
R. DeMaat,^{au} M. Demarteau,^{au} R. Demina,^{bn} P. Demine,^q D. Denisov,^{au} S.P. Denisov,^{aj}
S. Desai,^{bo} H.T. Diehl,^{au} M. Diesburg,^{au} M. Doets,^{ae} M. Doidge,^{am} H. Dong,^{bo}
S. Doulas,^{bg} L. Duflot,^o S.R. Dugad,^{aa} A. Duperrin,ⁿ O. Dvornikov,^{ag,cb} J. Dyer,^{bi}
A. Dyshkant,^{aw} M. Eads,^{aw} D. Edmunds,^{bi} T. Edwards,^{ao} J. Ellison,^{as} J. Elmsheuser,^x

J.T. Eltzroth,^{bt} V.D. Elvira,^{au} S. Eno,^{be} P. Ermolov,^{ai} O.V. Eroshin,^{aj} J. Estrada,^{au}
 D. Evans,^{an} H. Evans,^{bm} A. Evdokimov,^{ah} V.N. Evdokimov,^{aj} J. Fagan,^{au} J. Fast,^{au}
 F. Filthaut,^{af} W. Fisher,^{bl} H.E. Fisk,^{au} S.N. Fatakia,^{bf} L. Feligioni,^{bf} T. Ferbel,^{bn}
 M.J. Ferreira,^d F. Fiedler,^x T. Fitzpatrick,^{au} E. Flattum,^{au} R. Flores,^{au} J. Foglesong,^{au}
 M. Fortner,^{aw} H. Fox,^v C. Franklin,^{au} W. Freeman,^{au} S. Fu,^{au} S. Fuess,^{au}
 T. Gadfort,^{bx} C.F. Galea,^{af} E. Gallas,^{au} E. Galyaev,^{az} C. Garcia,^{bn} A. Garcia-Bellido,^{bx}
 J. Gardner,^{bb} V. Gavrilov,^{ah} P. Gay,^l D. Gelé,^r R. Gelhaus,^{as} K. Genser,^{au}
 C.E. Gerber,^{av} Y. Gershtein,^{bs} G. Ginther,^{bn} B. Gobbi,^{ax} T. Golling,^u V. Golovtsov,^{ak}
 B. Gómez,^g G. Gomez,^{be} R. Gomez,^g R. Goodwin,^{au} Y. Gornushkin,^{ag} K. Gounder,^{au}
 A. Goussiou,^{az} D. Graham,^{au} P.D. Grannis,^{bo} K. Gray,^{au} S. Greder,^r D.R. Green,^{au}
 J. Green,^{au} H. Greenlee,^{au} Z.D. Greenwood,^{bd} E.M. Gregores,^d Ph. Gris,^l J.-F. Grivaz,^o
 L. Groer,^{bm} S. Grünendahl,^{au} M.W. Grünewald,^{ab} W. Gu,^{au} J. Guglielmo,^{au} A. Gupta,^{aa}
 S.N. Gurzhiev,^{aj} G. Gutierrez,^{au} P. Gutierrez,^{br} A. Haas,^{bm} N.J. Hadley,^{be} E. Haggard,^{au}
 H.T. Haggerty,^{au} S. Hagopian,^{at} I. Hall,^{br} R.E. Hall,^{ar} C. Han,^{bh} L. Han,^{ao}
 K. Hanagaki,^{au} P. Hanlet,^{bg} S. Hansen,^{au} K. Harder,^{bc} R. Harrington,^{bg} J.M. Hauptman,^{ba}
 R. Hauser,^{bi} J. Hays,^{ax} E. Hazen,^{bf} T. Hebbeker,^t D. Hedin,^{aw} J.M. Heinmiller,^{av}
 A.P. Heinson,^{as} U. Heintz,^{bf} C. Hensel,^{bb} G. Hesketh,^{bg} M.D. Hildreth,^{az} R. Hirosky,^{bw}
 J.D. Hobbs,^{bo} B. Hoeneisen,^k M. Hohlfeld,^w S.J. Hong,^{ac} R. Hooper,^{bs} P. Houben,^{ae}
 Y. Hu,^{bo} J. Huang,^{ay} D. Huffman,^{au} I. Iashvili,^{as} R. Illingworth,^{au} A.S. Ito,^{au}
 S. Jabeen,^{bb} Y. Jacquier,^o M. Jaffré,^o S. Jain,^{br} V. Jain,^{bp} K. Jakobs,^v
 R. Jayanti,^{ap} A. Jenkins,^{an} R. Jesik,^{an} K. Johns,^{ap} M. Johnson,^{au} A. Jonckheere,^{au}
 P. Jonsson,^{an} H. Jöstlein,^{au} N. Jouravlev,^{ag} M. Juarez,^{au} A. Juste,^{au} A.P. Kaan,^{ae}
 M.M. Kado,^{ag} D. Käfer,^t W. Kahl,^{bc} S. Kahn,^{bp} E. Kajfasz,ⁿ A.M. Kalinin,^{ag}
 J. Kalk,^{bi} S.D. Kalmani,^{aa} D. Karmanov,^{ai} J. Kasper,^{bf} D. Kau,^{at} R. Kehoe,^{bu}
 S. Kermiche,ⁿ S. Kesisoglou,^{bs} A. Khanov,^{bn} A. Kharchilava,^{az} Y.M. Kharzheev,^{ag}
 H. Kim,^{bt} K.H. Kim,^{ac} N. Kirsch,^{bg} B. Klima,^{au} M. Klute,^u J.M. Kohli,^z
 E.V. Komissarov,^{ag} M. Kopal,^{br} V.M. Korablev,^{aj} A. Kostritski,^{aj} J. Kotcher,^{bp}
 B. Kothari,^{bm} A. Koubarovsky,^{ai} A.V. Kozelov,^{aj} J. Kozminski,^{bi} N. Kravchuk,^{ag}
 K. Krempetz,^{au} J. Krider,^{au} M.R. Krishnaswamy,^{aa} S. Krzywdzinski,^{au} B. Kubinski,^{au}
 N. Kuchinsky,^{ag} S. Kuleshov,^{ah} Y. Kulik,^{au} S. Kunori,^{be} A. Kupco,^q T. Kurča,^s
 R. Kwarcianny,^{au} S. Lager,^{al} N. Lahrichi,^q G. Landsberg,^{bs} M. Larwill,^{au} P. Laurens,^{bi}
 B. Lavigne,^o J. Lazoflores,^{at} A.-C. Le Bihan,^r G. Le Meur,^o P. Lebrun,^s S.W. Lee,^{ac}
 W.M. Lee,^{at} A. Leflat,^{ai} F. Lehner,^{47,*} R. Leitner,^h C. Leonidopoulos,^{bm} P. Lewis,^{an}
 J. Li,^{bt} Q.Z. Li,^{au} J.G.R. Lima,^{aw} D. Lincoln,^{au} C. Lindenmeyer,^{au} S.L. Linn,^{at}
 J. Linnemann,^{bi} V.V. Lipaev,^{aj} R. Lipton,^{au} M. Litmaath,^{au} J. Lizarazo,^g L. Lobo,^{an}

A. Lobodenko,^{ak} M. Lokajicek,^j A. Lounis,^r H.J. Lubatti,^{bx} L. Lueking,^{au} M. Lynker,^{az}
 A.L. Lyon,^{au} E. Machado,^{bf} A.K.A. Maciel,^{aw} R.J. Madaras,^{aq} P. Mättig,^y
 A. Magerkurth,^{bh} A.-M. Magnan,^m M. Maity,^{bf} N. Makovec,^o P.K. Mal,^{aa} S. Malik,^{bd}
 V.L. Malyshev,^{ag} H.S. Mao,^f Y. Maravin,^{au} D. Markley,^{au} M. Markus,^{bg} M. Martens,^{au}
 M. Martin,^{aw} G. Martin-Chassard,^o S.E.K. Mattingly,^{bs} M. Matulik,^{au} A.A. Mayorov,^{aj}
 R. McCarthy,^{bo} R. McCroskey,^{ap} M. McKenna,^{au} D. Meder,^w H.L. Melanson,^{au}
 A. Melnitchouk,^{bj} D. Mendoza,^{au} Y.P. Merekov,^{ag} M. Merkin,^{ai} K.W. Merritt,^{au}
 A. Meyer,^t H. Miettinen,^{bv} D. Mihalcea,^{aw} V. Mikhailov,^{ag,ca} D. Miller,^{au} J. Mitrevski,^{bm}
 N. Mokhov,^{au} J. Molina,^c N.K. Mondal,^{aa} H.E. Montgomery,^{au} R.W. Moore,^e S. Moua,^{au}
 N. Mokhov,^{au} M. Mulders,^{au} Y.D. Mutafov,^{bo} P. Nagaraj,^{aa} E. Nagy,ⁿ M. Narain,^{bf}
 V.S. Narasimhan,^{aa} A. Narayanan,^{ap} N.A. Naumann,^{af} H.A. Neal,^{bh} J.P. Negret,^g
 S. Nelson,^{at} R.T. Neuenschwander,^d P. Neustroev,^{ak} C. Noeding,^v A. Nomerotski,^{au}
 S.F. Novaes,^d A. Nozdrin,^{ag} T. Nunnemann,^x A. Nurczyk,^{au} E. Nurse,^{ao} V. O'Dell,^{au}
 D.C. O'Neil,^e V. Oguri,^c D. Olis,^{au} N. Oliveira,^c J. Olsen,^{au} N. Oshima,^{au}
 B.O. Oshinowo,^{au} G.J. Otero y Garzón,^{av} P. Padley,^{bv} N. Parashar,^{bd} J. Park,^{ac}
 S.K. Park,^{ac} J. Parsons,^{bm} R. Partridge,^{bs} N. Parua,^{bo} A. Patwa,^{bp} P.M. Perea,^{as}
 E. Perez,^q O. Peters,^{ae} P. Pétroff,^o M. Petteni,^{an} L. Phaf,^{ae} R. Piegaiia,^a
 P.L.M. Podesta-Lerma,^{ad} V.M. Podstavkov,^{au} Y. Pogorelov,^{az} P. Polosov,^{ah} B.G. Pope,^{bi}
 S. Porokhovoy,^{ag} W.L. Prado da Silva,^c W. Pritchard,^{au} I. Prokhorov,^{ag} H.B. Prosper,^{at}
 S. Protopopescu,^{bp} M.B. Przybycien,^{50,†} J. Qian,^{bh} A. Quadt,^u B. Quinn,^{bj}
 E. Ramberg,^{au} R. Ramirez-Gomez,^{bt} K.J. Rani,^{aa} M.V.S. Rao,^{aa} P.A. Rapidis,^{au}
 S. Rapisarda,^{au} J. Raskowski,^{aw} P.N. Ratoff,^{am} R.E. Ray,^{au} N.W. Reay,^{bc}
 R. Rechenmacher,^{au} L.V. Reddy,^{aa} T. Regan,^{au} S. Reucroft,^{bg} M. Rijssenbeek,^{bo}
 I. Ripp-Baudot,^r F. Rizatdinova,^{bc} M. Roco,^{au} C. Rotolo,^{au} C. Royon,^q P. Rubinov,^{au}
 R. Ruchti,^{az} R. Rucinski,^{au} N. Russakovich,^{ag} P. Russo,^{au} B. Sabirov,^{ag} G. Sajot,^m
 A. Sánchez-Hernández,^{ad} M.P. Sanders,^{ao} A. Santoro,^c B. Satyanarayana,^{aa} G. Savage,^{au}
 L. Sawyer,^{bd} T. Scanlon,^{an} R.D. Schamberger,^{bo} Y. Scheglov,^{ak} H. Schellman,^{ax}
 P. Schieferdecker,^x C. Schmitt,^y A.A. Schukin,^{aj} A. Schwartzman,^{bl} R. Schwienhorst,^{bi}
 S. Sengupta,^{at} H. Severini,^{br} E. Shabalina,^{av} M. Shamim,^{bc} H.C. Shankar,^{aa} V. Shary,^q
 P. Sheahan,^{au} W.D. Shephard,^{az} A.A. Shishkin,^{ag} D. Shpakov,^{bg} M. Shupe,^{ap}
 R.A. Sidwell,^{bc} V. Simak,ⁱ V. Sirotenko,^{au} P. Skubic,^{br} P. Slattery,^{bn} G. Smith,^{au}
 R.P. Smith,^{au} K. Smolek,ⁱ G.R. Snow,^{bk} J. Snow,^{bq} S. Snyder,^{bp} S. Söldner-Rembold,^{ao}
 X. Song,^{aw} Y. Song,^{bt} L. Sonnenschein,^{bf} A. Sopczak,^{am} M. Sosebee,^{bt} K. Soustruznik,^h
 M. Souza,^b N. Spartana,^{au} B. Spurlock,^{bt} N.R. Stanton,^{bc} J. Stark,^m J. Steele,^{bd}
 A. Stefanik,^{au} J. Steinberg,^{ap} G. Steinbrück,^{bm} K. Stevenson,^{ay} V. Stolin,^{ah} A. Stone,^{av}

D.A. Stoyanova,^{aj} J. Strandberg,^{al} M.A. Strang,^{bt} M. Strauss,^{br} R. Ströhmer,^x
M. Strovink,^{aq} L. Stutte,^{au} S. Sumowidagdo,^{at} A. Sznajder,^c M. Talby,ⁿ
S. Tentindo-Repond,^{at} P. Tamburello,^{ap} W. Taylor,^e P. Telford,^{ao} J. Temple,^{ap}
N. Terentyev,^{ak} V. Teterin,^{ag} E. Thomas,ⁿ J. Thompson,^{au} B. Thooris,^q
D. Toback,^{be} V.V. Tokmenin,^{ag} C. Tolian,^{au} M. Tomoto,^{au} D. Tompkins,^{ap} T. Toole,^{be}
J. Torborg,^{az} F. Touze,^o S. Towers,^{bo} T. Trefzger,^w S. Trincas-Duvold,^p B. Tuchming,^q
C. Tully,^{bl} A.S. Turcot,^{bp} P.M. Tuts,^{bm} M. Utes,^{au} L. Uvarov,^{ak} S. Uvarov,^{ak}
S. Uzunyan,^{aw} B. Vachon,^e P. van Gemmeren,^{au} R. Van Kooten,^{ay} W.M. van Leeuwen,^{ae}
N. Varelas,^{av} E.W. Varnes,^{ap} I.A. Vasilyev,^{aj} M. Vaupel,^y M. Vaz,^d P. Verdier,^o
L.S. Vertogradov,^{ag} M. Verzocchi,^{be} M. Vigneault,^{az} F. Villeneuve-Seguiet,^{an}
P.R. Vishwanath,^{aa} J.-R. Vlimant,^p E. Von Toerne,^{bc} A. Vorobyov,^{ak} M. Vreeswijk,^{ae}
T. Vu Anh,^o V. Vysotsky,^{ag,cc} H.D. Wahl,^{at} R. Walker,^{an} L. Wang,^{be} Z.-M. Wang,^{bo}
J. Warchol,^{az} M. Warsinsky,^u G. Watts,^{bx} M. Wayne,^{az} M. Weber,^{au} H. Weerts,^{bi}
M. Wegner,^t N. Wermes,^u A. White,^{bt} V. White,^{au} D. Whiteson,^{aq} D. Wicke,^{au}
T. Wijnen,^{af} D.A. Wijngaarden,^{af} N. Wilcer,^{af} H. Willutzki,^{bp} G.W. Wilson,^{bb}
S.J. Wimpenny,^{as} J. Wittlin,^{bf} M. Wobisch,^{au} J. Womersley,^{au} D.R. Wood,^{bg}
T.R. Wyatt,^{ao} Q. Xu,^{bh} N. Xuan,^{az} R. Yamada,^{au} M. Yan,^{be} R. Yarema,^{au} T. Yasuda,^{au}
sY.A. Yatsunenko,^{ag} Y. Yen,^y K. Yip,^{bp} F. Yoffe,^{au} S.W. Youn,^{ax}
J. Yu,^{bt} A. Yurkewicz,^{bi} A. Zabi,^o M. Zanabria,^{au} A. Zatserklyaniy,^{aw} M. Zdrazil,^{bo}
C. Zeitnitz,^w D. Zhang,^{au} X. Zhang,^{br} T. Zhao,^{bx} Z. Zhao,^{bh} B. Zhou,^{bh}
J. Zhu,^{be} M. Zielinski,^{bn} D. Zieminska,^{ay} A. Zieminski,^{ay} R. Zitoun,^{bo} T. Zmuda,^{au}
V. Zutshi,^{aw} S. Zviagintsev,^{aj} E.G. Zverev,^{ai} and A. Zylberstein^q

^a*Universidad de Buenos Aires, Buenos Aires, Argentina*

^b*LAFEX, Centro Brasileiro de Pesquisas Físicas, Rio de Janeiro, Brazil*

^c*Universidade do Estado do Rio de Janeiro, Rio de Janeiro, Brazil*

^d*Instituto de Física Teórica, Universidade Estadual Paulista, São Paulo, Brazil*

^e*Simon Fraser University, Burnaby, Canada, University of Alberta, Edmonton, Canada,*

McGill University, Montreal, Canada and York University, Toronto, Canada

^f*Institute of High Energy Physics, Beijing, People's Republic of China*

^g*Universidad de los Andes, Bogotá, Colombia*

^h*Charles University, Center for Particle Physics, Prague, Czech Republic*

ⁱ*Czech Technical University, Prague, Czech Republic*

^j*Institute of Physics, Academy of Sciences, Center for Particle Physics, Prague, Czech Republic*

^k*Universidad San Francisco de Quito, Quito, Ecuador*

^l*Laboratoire de Physique Corpusculaire, IN2P3-CNRS, Université Blaise Pascal, Clermont-Ferrand, France*

^m*Laboratoire de Physique Subatomique et de Cosmologie, IN2P3-CNRS, Université de Grenoble 1, Grenoble, France*

ⁿ*CPPM, IN2P3-CNRS, Université de la Méditerranée, Marseille, France*

^o*Laboratoire de l'Accélérateur Linéaire, IN2P3-CNRS, Orsay, France*

^p*LPNHE, Universités Paris VI and VII, IN2P3-CNRS, Paris, France*

^q*DAPNIA/Service de Physique des Particules, CEA, Saclay, France*

^r*IReS, IN2P3-CNRS, Université Louis Pasteur, Strasbourg, France and Université de Haute Alsace, Mulhouse, France*

^s*Institut de Physique Nucléaire de Lyon, IN2P3-CNRS, Université Claude Bernard, Villeurbanne, France*

^t*RWTH Aachen, III. Physikalisches Institut A, Aachen, Germany*

^u*Universität Bonn, Physikalisches Institut, Bonn, Germany*

^v*Universität Freiburg, Physikalisches Institut, Freiburg, Germany*

^w*Universität Mainz, Institut für Physik, Mainz, Germany*

^x*Ludwig-Maximilians-Universität München, München, Germany*

^y*Fachbereich Physik, University of Wuppertal, Wuppertal, Germany*

^z*Panjab University, Chandigarh, India*

^{aa}*Tata Institute of Fundamental Research, Mumbai, India*

^{ab}*University College Dublin, Dublin, Ireland*

^{ac}*Korea Detector Laboratory, Korea University, Seoul, Korea*

^{ad}*CINVESTAV, Mexico City, Mexico*

^{ae}*FOM-Institute NIKHEF and University of Amsterdam/NIKHEF, Amsterdam, The Netherlands*

^{af}*University of Nijmegen/NIKHEF, Nijmegen, The Netherlands*

^{ag}*Joint Institute for Nuclear Research, Dubna, Russia*

^{ah}*Institute for Theoretical and Experimental Physics, Moscow, Russia*

^{ai}*Moscow State University, Moscow, Russia*

^{aj}*Institute for High Energy Physics, Protvino, Russia*

^{ak}*Petersburg Nuclear Physics Institute, St. Petersburg, Russia*

^{al}*Lund University, Lund, Sweden, Royal Institute of Technology and Stockholm University, Stockholm, Sweden and Uppsala University, Uppsala, Sweden*

^{am}*Lancaster University, Lancaster, United Kingdom*

^{an}*Imperial College, London, United Kingdom*

^{ao}*University of Manchester, Manchester, United Kingdom*

^{ap}*University of Arizona, Tucson, Arizona 85721, USA*

^{aq}Lawrence Berkeley National Laboratory and University of California, Berkeley, California 94720, USA

^{ar}California State University, Fresno, California 93740, USA

^{as}University of California, Riverside, California 92521, USA

^{at}Florida State University, Tallahassee, Florida 32306, USA

^{au}Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA

^{av}University of Illinois at Chicago, Chicago, Illinois 60607, USA

^{aw}Northern Illinois University, DeKalb, Illinois 60115, USA

^{ax}Northwestern University, Evanston, Illinois 60208, USA

^{ay}Indiana University, Bloomington, Indiana 47405, USA

^{az}University of Notre Dame, Notre Dame, Indiana 46556, USA

^{ba}Iowa State University, Ames, Iowa 50011, USA

^{bb}University of Kansas, Lawrence, Kansas 66045, USA

^{bc}Kansas State University, Manhattan, Kansas 66506, USA

^{bd}Louisiana Tech University, Ruston, Louisiana 71272, USA

^{be}University of Maryland, College Park, Maryland 20742, USA

^{bf}Boston University, Boston, Massachusetts 02215, USA

^{bg}Northeastern University, Boston, Massachusetts 02115, USA

^{bh}University of Michigan, Ann Arbor, Michigan 48109, USA

^{bi}Michigan State University, East Lansing, Michigan 48824, USA

^{bj}University of Mississippi, University, Mississippi 38677, USA

^{bk}University of Nebraska, Lincoln, Nebraska 68588, USA

^{bl}Princeton University, Princeton, New Jersey 08544, USA

^{bm}Columbia University, New York, New York 10027, USA

^{bn}University of Rochester, Rochester, New York 14627, USA

^{bo}State University of New York, Stony Brook, New York 11794, USA

^{bp}Brookhaven National Laboratory, Upton, New York 11973, USA

^{bq}Langston University, Langston, Oklahoma 73050, USA

^{br}University of Oklahoma, Norman, Oklahoma 73019, USA

^{bs}Brown University, Providence, Rhode Island 02912, USA

^{bt}University of Texas, Arlington, Texas 76019, USA

^{bu}Southern Methodist University, Dallas, Texas 75275, USA

^{bv}Rice University, Houston, Texas 77005, USA

^{bw}University of Virginia, Charlottesville, Virginia 22901, USA

^{bx}University of Washington, Seattle, Washington 98195, USA

^{ca}Visitor from Institute for Nuclear Problems of the Belarussian State University, Minsk, Belarus

^{cb} Visitor from National Scientific and Educational Center of Particle and High Energy Physics, Belarussian State University, Minsk, Belarus

^{cd} Visitor from Research and Production Corporation “Integral,” Minsk, Belarus

Abstract

The DØ experiment enjoyed a very successful data-collection run at the Fermilab Tevatron collider between 1992 and 1996. Since then, the detector has been upgraded to take advantage of improvements to the Tevatron and to enhance its physics capabilities. We describe the design and performance of the new elements of the detector, including the silicon microstrip tracker, central fiber tracker, solenoidal magnet, preshower detectors, forward muon detector, and forward proton detector. We also present the associated electronics, triggering, and data acquisition systems, along with the design and implementation of software specific to DØ.

1 Introduction

The DØ experiment was proposed in 1983 to study proton-antiproton collisions at $\sqrt{s} = 1.8$ TeV at the Fermilab Tevatron collider. The focus of the experiment was the study of high mass states and large p_T phenomena. The detector performed very well during Run I of the Tevatron, 1992–1996, leading to the discovery of the top quark [1] and measurement of its mass [2–6], a precision measurement of the mass of the W boson [7–13], detailed analysis of gauge boson couplings [14–22], studies of jet production [23–26], and greatly improved limits on the production of new phenomena such as leptoquarks [27–33] and supersymmetric particles [34–41], among many other accomplishments [42].

During Run I, the Tevatron operated using six bunches each of protons and antiprotons, with 3500 ns between bunch crossings and a center-of-mass energy of 1.8 TeV. The peak luminosity was typically $1\text{--}2 \times 10^{31} \text{ cm}^{-2}\text{s}^{-1}$ and approximately 120 pb^{−1} of data were recorded by DØ. Following the completion of the new Main Injector and associated Tevatron upgrades, the collider began running again in 2001. In Run II, the Tevatron is operated with 36 bunches of protons and antiprotons with a bunch spacing of 396 ns, and at an increased center-of-mass energy of 1.96 TeV. The instantaneous luminosity will increase by at least a factor of ten to greater than $10^{32} \text{ cm}^{-2}\text{s}^{-1}$ and more than 4 fb^{−1} of data are expected to be recorded.

To take advantage of these improvements in the Tevatron and to enhance the physics reach of the experiment, we have significantly upgraded the DØ

detector. The detector consists of three major subsystems: central tracking detectors, uranium/liquid-argon calorimeters, and a muon spectrometer. The original DØ detector is described in detail in Ref. [43]. The central tracking system has been completely replaced; the old system lacked a magnetic field and suffered from radiation damage, and improved tracking technologies are now available. The new system includes a silicon microstrip tracker and a scintillating-fiber tracker located within a 2 T solenoidal magnet. The silicon microstrip tracker is able to identify displaced vertices for b -quark tagging; the magnetic field provides the opportunity for E/p measurements for electron identification and calorimeter calibration, opens new capabilities for τ identification and hadron spectroscopy, and allows precision muon momentum measurement. Between the solenoidal magnet and the central calorimeter and in front of the forward calorimeters, preshower detectors have been added for improved electron identification. The forward muon system's proportional drift chambers have been replaced by mini drift tubes and trigger scintillation counters that can withstand the harsh radiation environment and additional shielding has been added. In the central region, scintillation counters have been added for improved muon triggering. We have also added a forward proton detector for the study of diffractive physics. A side view of the upgraded DØ detector is shown in Figure 1.

The large reduction in the bunch spacing required the improvement of the read-out electronics and the implementation of pipelining for the front-end signals from the tracking, calorimeter, and muon systems. The calorimeter preamplifiers and signal-shaping electronics have been replaced, as have all of the electronics for the muon system. A new trigger level has been added to cope with the higher collision rate, and trigger hardware has been installed to identify displaced secondary vertices for b -quark tagging. Muon triggering has been enhanced by the addition of scintillation counters in the central and forward regions.

A significant improvement to the detector resulted from the removal of the old Main Ring beam pipe from the calorimeters. During Run I, the Main Ring was used to accelerate protons for antiproton production while the Tevatron operated in collider mode. Losses from the Main Ring produced spurious energy deposits in the calorimeters and muon system, and most triggers were not accepted while Main Ring protons passed through the detector. Removal of the Main Ring increased the livetime of the detector by approximately 10%, depending on the trigger.

In the following sections of this paper, we describe the design and performance of the upgraded DØ detector. The new central tracking system and solenoidal magnet are presented in Sections 2 and 3, respectively. The preshower detectors are described in Section 4. The calorimeters are briefly described in Section 5 along with the new calorimeter electronics. The muon system is dis-




Fig. 1. Diagram of the upgraded DØ detector, as installed in the collision hall and viewed from inside the Tevatron ring. The $+z$ axis is to the right, $+y$ is up, and $+x$ is out of the page.

cussed in Section 6. The new forward proton detector is presented in Section 7. The triggering and data acquisition systems are described in Sections 9 and 10. Detector controls and monitoring and an overview of the software components of the experiment are covered in Sections 11 and 12.

In detector description and data analysis, we use a right-handed coordinate system in which the z -axis is along the proton direction and the y -axis is upward (Figure 1). The angles ϕ and θ are the azimuthal and polar angles, respectively. The r coordinate denotes the perpendicular distance from the z axis. The pseudorapidity, $\eta = -\ln[\tan(\theta/2)]$, approximates the true rapidity, $y = 1/2 \ln[(E+p_z c)/(E-p_z c)]$, for finite angles in the limit that $(mc^2/E) \rightarrow 0$. We use the term “forward” to describe the regions at large $|\eta|$. Since some particles escape down the beam pipe and the initial momenta of the interacting partons are not known, we measure momentum, energy and missing energy in the direction perpendicular to the beam axis — i.e. transverse momentum (p_T), transverse energy (E_T), and missing transverse energy (\cancel{E}_T).

2 Central tracking

Excellent tracking in the central region is necessary for our studies of top quark, electroweak, and b physics and to search for new phenomena, including the Higgs boson. The central tracking system consists of the silicon microstrip tracker (SMT) and the central fiber tracker (CFT) surrounded by a solenoidal magnet. Working together, the two detectors locate the primary interaction vertex with a resolution of about $35\ \mu\text{m}$ along the beamline. They can tag b -quark jets with an impact parameter resolution less than $15\ \mu\text{m}$ in $r - \phi$ for particles with $p_T > 10\ \text{GeV}/c$ at $|\eta| = 0$. The high resolution of the vertex position allows good measurement of lepton p_T , jet E_T , and \cancel{E}_T . Calibration of the electromagnetic calorimeter using the E/p ratio of electrons is now possible.

Both the SMT and CFT provide tracking information for the trigger. The SMT provides digitized analog signals to the Level 2 and 3 trigger systems and is used to trigger on displaced vertices from b -quark decay. The CFT provides a fast and continuous readout of discriminator signals to the Level 1 trigger system; information from these discriminator signals is sent to Level 2 upon a Level 1 trigger accept. The Level 3 trigger receives a slower readout of the CFT's digitized analog signals, in addition to the discriminator information available at Level 1 and Level 2.

A schematic view of the central tracking system is shown in Figure 2.




Fig. 2. The new central tracking system. Also shown are the locations of the solenoid, the preshower detectors, luminosity monitor, and the calorimeters. Not to scale.

2.1 Silicon microstrip tracker

The SMT provides both tracking and vertexing over nearly the full η coverage of the calorimeter and muon systems. Design of the detector, electronics, and cooling are, in large part, dictated by the accelerator environment. The length of the interaction region ($\sigma \approx 25$ cm) sets the length scale of the device. With a long interaction region, it is difficult to deploy detectors such that the tracks are generally perpendicular to detector surfaces for all η . This led us to a design of barrel modules interspersed with disks in the center and assemblies of disks in the forward regions. The barrel detectors primarily measure the $r - \phi$ coordinate and the disk detectors measure $r - z$ as well as $r - \phi$. Thus vertices for high η particles are reconstructed in three dimensions by the disks, and vertices of particles at small values of η are measured in the barrels and central fiber tracker. This design poses difficult mechanical challenges in the arrangement of the detector types to provide space for cooling and cables while minimizing dead area.




Fig. 3. The disk/barrel design of the silicon microstrip tracker.

An isometric view of the SMT is shown in Figure 3. The detector has six barrels in the central region. Each barrel has four silicon readout layers. The silicon modules installed in the barrels are called “ladders.” Layers 1 and 2 have twelve ladders each; layers 3 and 4 have twenty-four ladders each, for a total of 432 ladders. Each barrel is capped at high $|z|$ with a disk of twelve double-sided wedge detectors, called an “F-disk.” Forward of each disk/barrel assembly is a unit consisting of three F-disks. In the far forward regions, two large-diameter disks, “H-disks,” provide tracking at high $|\eta|$. Twenty-four full wedges, each consisting of two back-to-back single-sided “half” wedges, are mounted on each H-disk. There are 144 F-wedges and 96 full H-wedges in the tracker; each side of a wedge (upstream and downstream) is read out independently. There is a grand total of 912 readout modules, with 792,576 channels. The centers of the H-disks are located at $|z| = 100.4, 121.0$ cm; the F-disks are at $|z| = 12.5, 38.2, 43.1, 48.1, \text{ and } 53.1$ cm.

2.1.1 Sensors

2.1.1.1 Sensor types The detector uses a combination of single-sided (SS), double-sided (DS), and double-sided double-metal (DSDM) technologies. Silicon sensors were obtained from three manufacturers. All barrel sensors were produced by Micron Semiconductor LTD [44]. The wedges for the F-disks were produced by Micron Semiconductor LTD and Canberra Eurisys Mesures [45]. The wedge sensors for the H-disks were manufactured by ELMA [46]; these sensors use intermediate strips for charge interpolation. Single-sided and double-sided devices were produced from high resistivity 4” silicon wafers, with crystal orientation $\langle 111 \rangle$ and $\langle 100 \rangle$. The 90° stereo sensors used in layers 1 and 3 of the four centermost barrels are double-sided double-metal sensors, manufactured using $\langle 100 \rangle$ 6” wafers. Isolation on the n-side of all double-sided sensors is provided by p-stop implants. All traces are biased using polysilicon resistors. Table 1 shows the sensor types used in the SMT and their locations.

Table 1

Characteristics and deployment of various sensor types in the SMT. i indicates the length of the inner H-disk sensor; o is the length of the outer H-disk sensor.

Module	Type	Layer	Pitch (μm) p/n	Length (cm)	Inner radius (cm)	Outer radius (cm)
F-disks	DS	–	50/62.5	7.93	2.57	9.96
H-disks	SS	–	40 80 readout	76.3 ⁱ 63.3 ^o	9.5	26
Central barrels (4)	DSDM	1, 3	50/153.5	12.0	2.715	7.582
	DS	2, 4	50/62.5	6.0	4.55	10.51
Outer barrels (2)	SS	1, 3	50	6.0	2.715	7.582
	DS	2, 4	50/62.5	6.0	4.55	10.51

Disk sensors are trapezoids with readout strips arranged parallel to the long edge of the devices. This provides an effective 30° stereo angle for the double-sided F-disks. A wedge for the H-disks consists of a pair of single-sided half-wedges mounted back-to-back, giving an effective stereo angle of 15° . This arrangement means that the strip length varies for strips which originate past the base of the trapezoid. There are three types of sensors in the central barrels. The second and fourth layers use double-sided stereo sensors with the n-side implants at a 2° angle with respect to the p-side axial strips. Two of these sensors are bonded together to form a 12-cm ladder. The first and third layers of the outer barrels use single-sided sensors with axial strips and, again, two sensors are bonded together to make one 12-cm ladder. The inner four barrels use single 12-cm-long 90° stereo sensors. Ninety degree readout was achieved by using a second metal layer on the n-side insulated from the first metal by $3\ \mu\text{m}$ of PECVD (plasma enhanced chemical vapor deposition) silicon oxide. Two readout strips on this side are multiplexed to a single readout channel. Implants on the n-side are isolated by individual p-stop frames in addition to a common p-stop enclosure.

2.1.1.2 Sensor testing All sensors were required to exhibit leakage current less than $260\ \text{nA}/\text{cm}^2$ and each channel was tested for leakage current, AC coupling capacitance, and AC coupling leakage to 80 V. The AC coupling leakage test is especially important for the double-sided sensors since these capacitors are required to stand off the depletion voltage. Several classes of problems were found during the testing:

- Areas of low interstrip resistance — Washing bad detectors in pure water at the vendor sometimes cured this problem.
- Isolation-implant shorts — These occur in all detector types but the effects




Fig. 4. Depletion voltage as a function of fluence as measured for SMT modules in the Fermilab Booster.

are particularly severe in the DSDM detectors where a single isolation short can affect a region of approximately twenty channels. Strip-by-strip DC scans were performed at the manufacturer to identify problems early in processing.

- Poorly controlled bias resistance — Sensors with bias resistors as high as $10\text{ M}\Omega$ ($5\times$ nominal) were eventually accepted.
- Microdischarge breakdown — This breakdown was observed with modest negative voltage applied to the p-side of double-sided sensors. The onset voltage of breakdown was measured for each sensor and used to set the split between p(-)- and n(+)-side bias voltage.

2.1.1.3 Radiation damage studies A set of radiation damage qualification studies was performed at the Fermilab Booster and the Radiation Laboratory at the University of Massachusetts Lowell. In general, the results of these studies conformed to behavior expected from RD48 [47] parameterizations. The exception was the depletion voltage (V_d) behavior of the DSDM devices, shown in Figure 4. These sensors exhibit a rise in depletion voltage almost twice as fast as that of the single-sided detectors. To determine if the anomalous behavior is due to the bulk silicon properties, photodiode test structures from the same wafer were irradiated in the booster. These all showed normal behavior. It is likely that the rapid rise in depletion voltage is related to the PECVD inter-metal isolation layer.

Microdischarge breakdown is caused by large fields near the junction implant. This high-field region is expected to move with the junction from the p-side to the n-side upon type inversion. In addition, fixed positive charge in the oxide insulation layers will tend to increase the field on the p-side and reduce it on the n-side. This is indeed what we observe. Breakdown moves from the p-side to the n-side after type inversion and the threshold increases to approximately 150 V after about 2 Mrad. Microdischarge in the DSDM detectors will limit the lifetime of the SMT to an integrated luminosity of $3.5\text{--}6\text{ fb}^{-1}$.




Fig. 5. Double-sided ladder design, n-side. The SVX chips shown as dashed lines are located on the p-side of the ladder.

2.1.2 Mechanical design

2.1.2.1 Ladders Sensor ladder assemblies were designed for low mass, precise alignment, and good thermal performance. With the exception of the DSDM ladders, which use a single 12-cm sensor, ladders were constructed using two 6-cm sensors. All ladder assembly was done under the control of a coordinate measuring machine (CMM). The two sensors are aligned using the fiducials on the sensors. Alignment was achieved during assembly using locating notches in precisely machined beryllium pieces on both the “active” (SVXIIe-readout-chip carrying, see Section 2.1.3) and “passive” sides of the ladder. These notches correspond to posts on the support bulkheads and are aligned to fiducials on the sensor. Carbon-boron fiber/Rohacell rails bridge the beryllium pieces and maintain the flatness of the ladder assemblies.

A sketch of a double-sided 2° ladder with nine SVXIIe readout chips is shown in Figure 5. Figure 6 shows a ladder on an assembly fixture with the flex hybrid unfolded, spring loaded “pushers,” and support rails. Each ladder and wedge was surveyed using an automatic optical CMM after assembly.

2.1.2.2 Wedges In contrast to the barrel assemblies, F- and H-disks are planar modules, allowing for simple optical alignment during assembly. A major design constraint was to provide minimum dead space between the disk/barrel assemblies. F-wedges use double-sided sensors with separate 8- and 6-chip high density interconnects (HDIs, see below) for each side. The larger strip pitch ($62.5\ \mu\text{m}$ vs $50\ \mu\text{m}$) on the n-side requires an adapter to




Fig. 6. Double-sided ladder during assembly

match the SVXIIe readout pitch. The $50\text{ }\mu\text{m}$ -thick copper-clad Kapton [48] pitch adaptors are part of the cooling path to the p-side and allow the six n-side SVXIIe chips to be located outboard of the cooling channel. The double-sided hybrid assembly required a complex 12-layer vacuum lamination with pitch adaptor, Tedlar [49], Ablefilm adhesive [50], Kapton [48], beryllium, and HDI layers.

H-disk wedges are made from back-to-back single-sided sensors. In this case, the precise front-to-back alignment was provided in the wedge assembly process. A special dual-camera assembly fixture was developed to simultaneously image fiducials on the top and bottom sensors. This allowed front-to-back alignment of the two-sensor wedges of better than $15\text{ }\mu\text{m}$. An aluminized silicon pitch adapter, fabricated on the same wafer as the sensor, was used to match the $80\text{ }\mu\text{m}$ readout pitch to the SVXIIe input.

2.1.2.3 Supports and assembly Barrel ladders are supported by beryllium bulkheads machined with posts and pinholes for ladder support. Each barrel has a thick “active” bulkhead containing the cooling channels and connections to the outer support cylinder, and a thin “passive” bulkhead that serves to set the spacing of the ends of the ladders without readout chips. Ladders were lowered onto the posts using a special insertion fixture and fixed into place with pins. Thermal contact to the cooled bulkhead ledge is made with thermal grease. Each ladder is grounded to the bulkheads with conducting epoxy. Edges of the beryllium pieces were measured with the touch probe of the assembly CMM and these measurements were correlated to the beryllium and sensor fiducial mark measurements taken with the optical CMM to provide the final ladder position. Cables (HDI “tails”) are routed between bar-




Fig. 7. Cross section of the SMT disk/barrel module showing ladders mounted on the beryllium bulkhead, sample cable paths, three of twelve F-disk wedges, carbon fiber support structure, and the low mass cable stack.

rel sublayers inboard of ladders, so no inter-module space is taken by cables. The tails are trimmed to length and coupled to a “card edge” style Hirose connector [51] on the low mass cables on the outer surface of the support structure.

Disks are supported by beryllium rings. Wedges are located on alternate sides of the ring with sufficient overlap to eliminate dead regions. Wedges were manually aligned under a CMM and secured with screws. Finished F-disks were then assembled into a disk/barrel central module or one of the three-disk modules at the end of the central disk/barrel section.

Overall support of the SMT (exclusive of the H-disks) is provided by two double-walled carbon fiber cylinders with walls spaced by carbon fiber ribs to eliminate differential contraction. North and south half-cylinders are independent structures. This limits the size of the units, allowing installation of the SMT in the limited space available in the collision hall. The central upper section of each half-cylinder is removed for placement of the disk/barrel modules. Each module is supported by adjustable kinematic mounts. Cables and services are accessed through holes in the cylinder whose outer surface is used for routing the low mass cables and water manifolds. Final alignment is provided by sapphire balls mounted on the bulkheads, which are accessed with touch probes through additional holes in the support cylinder. The disk/barrel half cylinders are supported from the inner central fiber tracker barrel using mounts glued into place. Figure 7 shows the disk/barrel module within its support cylinder. H-disks were installed on separate mounts suspended from CFT layer 3.

2.1.2.4 Cooling Bulk damage to the silicon, leading to type inversion and increasing depletion voltage, limits the maximum operating temperature of the detector. Heat generated by the SVXIIe chips must flow through the Kapton HDI flex circuit, epoxy, beryllium heat spreaders and, for the top-side chips, through the silicon and the bottom-side readout structure.

Coolant is a water/30% ethylene glycol mixture supplied at -10° C by two redundant 4.4 kW chillers. The coolant flows through channels machined in the active beryllium bulkheads and ring supports and is maintained below atmospheric pressure to minimize the effect of any possible leaks. Pressure drop across detector elements is below 5 psi with a temperature rise in the channel of 1–1.5° C. A manifold is used for each side of the half-cylinder with flows through the modules controlled by restricting apertures. The hottest point of the silicon at the tip of the ladder is typically 10–15° C above the coolant temperature, maintaining the silicon temperature below 5° C.

2.1.3 Electronics

2.1.3.1 SVXIIe chip and hybrids The SMT is read out using the 128-channel SVXIIe chip [52]. The chip includes preamp, analog delay, digitization, and data sparsification. Input charge is integrated on the preamplifier for a train of beam crossings (typically twelve) and is reset during inter-bunch gaps. This charge is delivered to a 32-cell analog pipeline. Upon a Level 1 trigger accept, double-correlated sampling is performed on the appropriate cells and this analog information is fed to a parallel set of Wilkinson ADCs. Digitization utilizes both edges of a 53 MHz clock, providing 8 bits of analog information in 2.4 μ s. Readout is half as fast. Typical noise performance for a rise time setting of 100 ns is $490e + 50e/pF$. Fabrication was done using the 1.2 μ m UTMC radiation-hard process [53] with yields of approximately 60%.

Chips were mounted on Kapton flex circuits laminated to beryllium heat spreaders. Eight different types of flex were needed to accommodate the various detector and readout geometries. A readout cable “tail” is part of each HDI flex circuit. In the case of double-sided ladder detectors, a single circuit contains the readout for both the p- and n-sides with the Kapton folded over to sandwich the detector. The flex circuits are two-layer 50 μ m Kapton with 125 μ m line spacing and 50 μ m plated-through holes.

2.1.3.2 Readout Figure 8 shows the SMT readout chain. Trigger information is received via the Serial Command Link (SCL) by the sequencer crate controller. The SVXIIe sequencer provides timing and control signals for the SVXIIe chips on eight HDIs. These signals are regenerated by interface boards located on the side of the central calorimeter cryostat. The interface boards




Fig. 8. Diagram of the major components of the SMT readout system.

also control power and bias for the SVXIIe chips and provide interfaces to the monitoring systems and individual HDI circuit temperature and current trips.

Data from the HDIs are sent from the sequencers to VME readout buffer (VRB) memories via optical link fibers. The VRBC (VRB controller) receives trigger data from the SCL and uses that information to control the operation of the VRBs. A single-board computer is resident in the readout crate to collect and process detailed diagnostic information. Downloads and slow controls are provided by a MIL-1553 control system.

A large scale “10%” test was organized to test major detector components with pre-production versions of final readout components. This test was crucial for debugging readout hardware, testing termination schemes, and adjusting sequencer timings. Bugs were found in the SVXIIe chip that caused pedestal jumps and readout errors. Additional initialization states were added to the SVXIIe control sequence to solve these problems. An assembled barrel and F- and H-disks were tested in the 10% test to check system integration and search for noise problems. Cosmic ray data were taken with the barrel as a final check of system performance.

2.1.3.3 Cabling Clock, power, and signal quality and timing are critical to proper operation of the SVXIIe chip. The SMT is read out using low mass Kapton flex cables within the detector volume followed by high mass 80- and 50-conductor “pleated-foil” cables on the outside. The cables carry both power (except the 50-conductor high-mass cables) and digital signals. A pair of coaxial cables carries the differential clock. Low mass cables are routed along the half-cylinder and coupled to the 80-conductor pleated-foil cables at a ring of adapter cards located between the calorimeter cryostats.

2.1.4 Production and testing

Production of ladders and wedges and assembly into barrels and disks were performed at the Silicon Detector Facility (SiDet) at Fermilab. Figure 9 shows the fraction of completed 3-, 6-, and 9-chip wedges and F-disks as a function of time. Components such as chips, HDIs, and cables were tested before shipment to Fermilab. A number of sensor- and HDI-related problems were found which necessitated extensive quality control testing. Tests included:

- HDI burn-in — Each HDI was burned-in at room temperature for 24 hours. Current consumption, pedestal, and gain were monitored.
- Ladder/wedge debugging — Pedestal and noise were measured for each assembled module. Pinholes (shorted AC coupling capacitors) generated in the wirebonding process were identified in 0.8% of the channels and the associated wirebonds pulled. The onset voltage for microdischarge was determined on the n-side of double-sided detectors. Resistance between the HDI ground and beryllium was measured. Grounds were remade if the resistance was greater than 10 Ω .
- Ladder/wedge burn-in — Assembled and tested ladders and wedges were burned in at 5° C for 72 hours. Currents, pedestals, gain, data integrity, and noise were monitored.
- Laser test — Modules were tested using a 1064 nm laser on a computer-controlled $x - y$ table. This test allowed us to measure depletion voltage and tabulate broken and noisy channels.
- Assembly test — Noise and readout were tested for each ladder (wedge) as it was installed in a barrel (disk). Coherent noise was evident in badly grounded ladders and grounds were remade.
- Final cabling test — Readout was verified after low mass cables were attached. Upon leaving SiDet, 99.5% of the detectors were functional.

This extensive testing was made possible by the development of a “Stand Alone Sequencer” (SASEQ), a 3U VME module which could be read out simply using a PC-VME interface and Microsoft EXCEL. This allowed for easy replication of testing systems at Fermilab and universities.




Fig. 9. Fraction of completed 3-, 6-, and 9-chip wedges and F-disks as a function of time.

2.1.5 Operation

2.1.5.1 Operating experience The SMT has been included in physics data taking since the start of Run II in March 2001. Full electronics debugging was completed in October 2001, when 94% of the HDIs were functional. As of February 2005, 89% of HDIs were functional. Most operational problems have been peripheral to the silicon detector itself. These include latchup of op amps on the interface boards, low voltage power supply failures, and high leakage currents in high voltage distribution boxes.

The most serious detector problem is “grassy noise,” shown in Figure 10. This noise is characterized by large charge spikes which cover 10–20 strips and occur in about 20% of the events for affected devices. Leakage currents typically rise to greater than $100 \mu\text{A}$ within one hour of turn-on at the beginning of a store. This noise is confined to the Micron-supplied F-disk detectors. Both the SMT and CFT observe pedestal shifts that depend on the phase of the beam crossing with respect to the SVXIIe reset pulse. These shifts are typically 1–2 counts (27 counts/MIP) in most ladders but can be as large as 10 counts in a few that presumably have poor ground connections. Digital voltage supply current in the SVXIIe chip rises steadily if the chip is not properly initialized or not read out steadily and can cause individual HDIs to trip off in 5–10 minutes. These currents are constantly displayed and a special pulser is turned on if there is an extended interruption in data taking.




Fig. 10. Example of the grassy noise seen in the Micron-supplied F-disk detectors.

2.1.5.2 Alignment and calibration Signal/noise performance varies with detector type from 12:1 to 18:1. Coherent noise is typically one-third of the random noise but varies due to crossing-dependent pedestal shifts. Gains vary among detector types with the n-sides 5–15% lower than the p-sides due to the larger load capacitance. Pulse height information from the SVXIIe is used to calculate cluster centroids and can also be used for dE/dx tagging of low momentum tracks. Figure 11 shows dE/dx distributions after corrections for gain and incident angle are made.

Detector alignment was transferred from optical CMM measurements of detector fiducials to ladder beryllium features to barrel sapphire balls to half-cylinder targets to the DØ coordinate system. Final alignment for the combined SMT-CFT tracking system is better than $10\ \mu\text{m}$ (Figure 12).




Fig. 11. Distribution of energy loss for a kaon-enriched sample of tracks showing π , K , and proton bands.




Fig. 12. Axial residual distribution using tracks with $p_T > 3$ GeV/ c for (a) unaligned and (b) aligned geometries for central barrel detectors. The residual is the distance between the SMT hit and the track. The track fit was done using all SMT and CFT hits except the SMT hit in question. The simulated resolution of the residual distribution for a perfectly aligned detector is about $16 \mu\text{m}$.

2.2 Central fiber tracker

The CFT consists of scintillating fibers mounted on eight concentric support cylinders and occupies the radial space from 20 to 52 cm from the center of the

beampipe. The two innermost cylinders are 1.66 m long; the outer six cylinders are 2.52 m long. Each cylinder supports one doublet layer of fibers oriented along the beam direction (z) and a second doublet layer at a stereo angle of $+3^\circ$ (u) or -3° (v). Doublet layers with fibers oriented along the beam axis are referred to as axial layers, while the doublet layers oriented at small angles are referred to as stereo layers. From the smallest cylinder outward, the fiber orientation is $zu - zv - zu - zv - zu - zv - zu - zv$. The scintillating fibers are coupled to clear fiber waveguides which carry the scintillation light to visible light photon counters (VLPCs, Section 2.2.4) for read out. The small fiber diameter gives the CFT an inherent doublet layer resolution of about $100\ \mu\text{m}$ as long as the location of the individual fibers is known to better than $50\ \mu\text{m}$. Scintillating fiber detectors are discussed in detail in Ref. [54].

Signals from the axial doublet layers are used to form a fast Level 1 hardware trigger based upon the number of track candidates above a specified p_T threshold (with a minimum threshold of $1.5\ \text{GeV}/c$). Level 1 track candidates are used for the Level 2 trigger, and the Level 3 trigger uses the full CFT readout.

2.2.1 *Fibers*

The scintillating fibers, including cladding, are $835\ \mu\text{m}$ in diameter and 1.66 or 2.52 m in length. They are optically connected to clear fiber waveguides of identical diameter which are 7.8 to 11.9 m long. The fibers were manufactured by Kuraray [55] and have a multi-clad structure consisting of a core surrounded by two claddings. The scintillating fiber is structurally and chemically similar to the clear fiber, but contains fluorescent dyes. The CFT uses about 200 km of scintillating fiber and 800 km of clear fiber.

Light production in the fibers is a multistep process. The base core material is polystyrene (PS). The PS is doped with the organic fluorescent dye paraterphenyl (pT) to about 1% by weight. Excitations in the PS are rapidly transferred to the pT via a non-radiative dipole-dipole interaction. pT has a rapid fluorescence decay (a few nanoseconds) and a short emission wavelength ($\approx 340\ \text{nm}$). The mean free path of the emitted light is only a few hundred microns in the PS. To get the light out of the detector, a secondary, or wave-shifter, dye, 3-hydroxyflavone (3HF), is added at a low concentration (1500 ppm). The 3HF is spectrally matched to the pT but has minimal optical self-absorption. The 3HF absorbs the 340 nm radiation from the pT and re-emits it at 530 nm which is well-transmitted in PS.

Surrounding the PS core, whose refractive index is $n = 1.59$, are two claddings, each approximately $25\ \mu\text{m}$ thick: an inner layer of polymethylmethacrylate (PMMA) with $n = 1.49$, and an outer layer of fluoro-acrylic with $n = 1.42$.




Fig. 13. Technique for curved scintillating fiber ribbon fabrication.

The PMMA inner cladding serves as a mechanical interface between the core and the outer cladding, which are mechanically incompatible. The multicladd fiber is both mechanically and optically superior to single-clad fiber and typical values of the attenuation length are about 5 m for the scintillating fiber and about 8 m for the clear fiber.

We observe the light from only one end of each scintillating fiber. The opposite end of each of the scintillating fibers was mirrored by sputtering with an aluminum coating that provides a reflectivity of 85 to 90%.

The scintillating fibers were made into ribbons consisting of 256 fibers in two layers of 128 fibers each. Precisely spaced grooves were machined into a long, 1/16"-thick piece of Delrin plastic. The spacing between the grooves varies between 928 and 993 μm and depends on the radius of the corresponding support cylinder. The grooved plastic was inserted into a rigid, curved backing with the desired radius and the scintillating fibers were laid in and glued together to form the doublet ribbons; the two layers of fiber are offset by one-half of the fiber spacing. The technique is illustrated in Figure 13 and creates curved ribbons to match the curvature of each support cylinder without machining precisely-spaced grooves into a curved surface. Details on fiber lengths and spacing are provided in Table 2.

The readout ends of the fibers were carefully positioned and adhesively bonded into v-groove connectors, and then the mass-terminated ribbon and connector were polished to facilitate high efficiency light transmission across the connector joint. A polished curved connector is shown in Figure 14. Each 256-fiber waveguide bundle terminates in a matching curved connector. The connectors for each doublet fiber layer are different since the connectors must have the correct curvature for each layer. The light transmission through the v-groove connectors, with optical grease between the fiber ends, is approximately 95%.

After the fiber ribbons were fabricated, a ^{57}Co x-ray source was used to verify the accuracy of the fiber placement, the responses of the fibers, and the transmission efficiency of the connectors. The position of each fiber within a ribbon was verified with an accuracy of better than 25 μm rms.

Table 2

Design parameters of the CFT; $u = +3^\circ$, $v = -3^\circ$

Layer	Radius (cm)	Fibers/layer	Fiber separation (μm)	Active length (m)
A	20.04	1280×2	982.4	1.66
Au	20.22	1280×2	990.3	1.66
B	24.93	1600×2	978.3	1.66
Bv	25.13	1600×2	985.1	1.66
C	29.87	1920×2	976.1	2.52
Cu	30.05	1920×2	980.9	2.52
D	34.77	2240×2	974.4	2.52
Dv	34.95	2240×2	979.3	2.52
E	39.66	2560×2	971.7	2.52
Eu	39.86	2560×2	976.3	2.52
F	44.56	2880×2	970.0	2.52
Fv	44.74	2880×2	974.3	2.52
G	49.49	3200×2	969.8	2.52
Gu	49.67	3200×2	973.3	2.52
H	51.97	3520×2	926.1	2.52
Hv	52.15	3520×2	927.8	2.52




Fig. 14. A CFT fiber ribbon mass terminated via a v-groove connector.

2.2.2 Mechanical support structure

The eight support cylinders are each double-walled with a 0.25"-thick core of Rohacell [56]. The walls are constructed from linear carbon fibers impregnated with about 40% resin. To minimize sagging of the cylinders, the carbon fibers were applied in layers in the following pattern: $0^\circ/60^\circ/-60^\circ/\text{core}/-60^\circ/60^\circ/0^\circ$ where the numbers are the angles in degrees of the carbon fibers with respect to the axis of the cylinder. Each carbon fiber layer is about 0.0025" thick; the total thickness of carbon per cylinder is 0.020". After fabrication, the outside diameters of the cylinders were measured and compared

to a perfect cylinder. The radial deviation for all cylinders is approximately $100\text{ }\mu\text{m}$ rms.

The requirements of the Level 1 hardware trigger constrain the design of the CFT, particularly in the placement of the scintillating fiber ribbons. Offline reconstruction of the CFT data can incorporate corrections for misalignment of fibers, but that is not possible in the hardware trigger; the axial fibers had to be placed on the support cylinders such that any skew away from parallel to the beam axis is less than $200\text{ }\mu\text{m}$ from end to end. Consequently, the individual fiber ribbons were precisely positioned and bonded to the outer surface of each cylinder prior to nesting the cylinders. Global placement precision was approximately $35\text{ }\mu\text{m}$ rms for both axial and stereo layers. The connectors for the axial ribbons are all on one end of the cylinder, while the connectors for the corresponding stereo ribbons are all on the opposite end of the cylinder.

Successive cylinders are interlocked together by thin carbon-fiber annular rings which connect the inner surface of one cylinder's end ring to a carbon-fiber ring mounted on the outer fiber surface of the cylinder immediately inside. The fiber connectors and connecting rings seal the fiber volume.

For tracks traversing the detector at normal incidence, the thickness of each cylinder can be broken down as follows: 0.28% of a radiation length for the scintillating fibers, 0.32% for the carbon fiber support cylinder, 0.13% for the glue used to make ribbons out of fibers, and 0.17% for the glue used to attach the ribbons to the support cylinders.

2.2.3 Waveguides

The clear fiber waveguides range in length from 7.8 to 11.9 m, and use the clear fibers described in Section 2.2.1. The waveguides generally contain 256 clear fibers inside a flexible protective plastic light shield. One end is mass terminated in a curved Delrin connector machined to mate to the corresponding connector at the end of the scintillating fiber ribbon. The mass-terminated curved connector end was polished using a diamond fly cutter. About 40 cm from the other end of the waveguide, the 256 fibers are separated in ϕ into two groups of 128 fibers. Each group is covered with a flexible plastic protective light shield. These fibers are individually routed into the appropriate locations in two rectangular connectors that are designed to mate to the VLPC cassettes (Section 2.2.4). These molded plastic rectangular connectors are composed of Noryl N190 [57] with Celogen RA foam [58] to minimize distortions. After the fibers were potted into the rectangular connectors, the rectangular connectors were polished and the routing of the fibers was verified.

The clear fiber waveguides are routed from the ends of the CFT through the small gaps between the central and end calorimeter cryostats to the VLPCs




Fig. 15. Routing of the clear waveguide fibers on the face of the central cryostat.

located about 6 m below the central calorimeter cryostat. These gaps also contain the forward preshower detectors (Section 4.3), their waveguide fibers, the central preshower detector's waveguides (Section 4.2), and the SMT readout cables. The narrowest region, between a forward preshower detector and the solenoidal magnet, is 1.5" wide. The waveguides follow complex paths along the central calorimeter cryostat to minimize their depth in z and provide the proper timing at the VLPCs. The waveguide routing is illustrated in Figure 15.

2.2.4 Visible light photon counter cassettes

The light generated by the passage of charged particles through the scintillating fibers of the CFT is converted into electrical signals by the VLPCs housed in the VLPC cassettes. VLPCs are impurity-band silicon avalanche photodetectors that operate at 9 K, and are capable of detecting single photons [59,60]. They provide fast response, excellent quantum efficiency ($\geq 75\%$), high gain (17,000 to 65,000), low gain dispersion, and the capability of functioning in a high background environment.

VLPCs are fabricated by growing a series of doped and undoped silicon layers on silicon substrate wafers. Individual wafers yield a maximum of 176 VLPC chips after dicing, and each chip contains a two by four array of 1-mm-diameter pixels. Each eight-pixel chip is soldered to an aluminum nitride substrate, and the outputs from individual pixels are wirebonded to individual contact pads on the substrate. Non-uniformities in the VLPC production process result in variations in VLPC characteristics such as gain, quantum efficiency, and thermal noise rates among and across VLPC wafers [61]. Due to these variations, the bias voltage at which the VLPCs operate at optimal signal-to-noise ratio varies between 6 and 8 volts. To reduce the complexity of the bias voltage and threshold implementations in the readout electronics, the VLPC chips are carefully sorted and assigned to specific cassettes to allow for optimal performance.

VLPC cassettes mounted in cryostat slots provide the mechanical support, optical alignment, and appropriate operating environment for the VLPCs. Figure 16 shows an outside view of a cassette with a readout board attached. Each VLPC cassette houses 128 VLPC chips, and thus provides 1024 individual pixels of light-sensitive detector. (Details on the specially-designed VLPC cryostats are available in Ref. [62].) Cassettes are approximately 88 cm tall, 48 cm wide, and 4.4 cm thick. Individual 0.965-mm-diameter fibers inside each cassette guide light from the clear fiber waveguides to individual VLPC pixels, and flex circuits provide paths for the electrical signals from the VLPCs to the preamplifiers on the analog front end boards (AFEs) that are mounted on the cassette body. In addition to preamplifiers, the AFEs also provide trigger discrimination, temperature control, and bias-voltage control electronics. Details on the AFEs are provided in Section 2.2.5.

Each VLPC cassette consists of a cassette body housing eight modules. A view of the internal structure of a partially assembled cassette is shown in Figure 17. Each module is composed of a 128-fiber bundle and a cold-end assembly. The cassettes are designed so that modules can be readily extracted to allow repair of the cold-end assemblies. The 128-fiber bundle terminates at the top end in a rectangular molded optical connector (referred to as the warm-end optical connector), and at the bottom (cold) end in sixteen groups of eight fibers. The




Fig. 16. VLPC cassette supporting AFE readout boards as viewed from the left side. The VLPC hybrids are located on the isotherms housed inside the copper cup shown at the bottom of the figure.




Fig. 17. VLPC cassette body with left side body panel and side panels removed to show four installed modules.

ends of the fibers are polished and each group of eight fibers is glued into a molded plug. The plugs (and mating boxes) are precision-molded parts made of carbon-fiber-loaded polyphenylene sulfide (PPS) plastic. A polyurethane feedthrough block is cast around the 128 fibers and the flex circuits to form the barrier between the warm and the cold ends of the cassette. The fibers accept light from the clear fiber waveguides which are connected to the warm-end optical connectors at the top of the cassette and pipe the light to the VLPCs mounted in the cold-end assemblies.

The cold-end assembly hangs from the PPS plugs and consists of sixteen 8-channel VLPC hybrid assemblies supported on an isotherm. The inset in the




Fig. 18. Cold-end assembly for one 128-channel module of a VLPC cassette.

lower left corner of Figure 16 shows a view of a cold-end assembly supported on a 128-fiber bundle, and Figure 18 shows an exploded view of the components of a cold-end assembly.

A VLPC hybrid assembly is composed of a PPS box adhesively bonded to the aluminum nitride substrate upon which a VLPC chip is mounted. The molded PPS box is precisely positioned (maintaining $25\ \mu\text{m}$ tolerances) so that the eight polished fibers in the PPS plug are aligned over the individual VLPC pixels when the plug and box are mated, thus achieving good light collection efficiency while minimizing optical crosstalk between neighboring channels. Connectors are clamped across the contact pads on the aluminum nitride substrates and mated to flex circuits which transmit single-ended electrical signals between the cold-end and the AFE. Each module contains two flex circuits, and these flexible printed circuits are realized on $5\text{-}\mu\text{m}$ -thick adhesiveless copper coatings on a $51\text{-}\mu\text{m}$ -thick polyamide base material. The individual traces

are $76\text{ }\mu\text{m}$ wide with minimum $89\text{ }\mu\text{m}$ spacings between circuit features. Each flex circuit is 41.1 cm long. The isotherm also supports a calibrated carbon resistor that serves as a temperature sensor and resistors used as heaters, as well as the required springs and fasteners. The temperature sensor and heaters are employed to control the temperature of the VLPCs to within 0.1 K.

The cassette body can be viewed as composed of cold-end and warm-end mechanical structures. The cold-end structure (that portion of the cassette that is inserted in the cryostat) comprises several sub-assemblies: the feedthrough assembly, the G-10 side walls, the heat intercept assemblies, and the cold-end copper cup. The cassette is mounted to the top plate of the cryostat at the feedthrough assembly. This assembly provides the gas-tight seal for the cold, ≤ 2 psig, stagnant helium gas volume within the cryostat. Side walls composed of G-10 support the top (or liquid nitrogen) intercept which serves to cut off the flow of heat from the warm end. Another set of G-10 side walls support the bottom (or liquid helium) intercept. The bottom intercept supports the copper cup which surrounds the isotherms at the ends of the modules.

The warm-end structure is made of parallel tin-plated aluminum panels separated by spacer bars that form a protective box for the optical fibers. Rails mounted on the aluminum panels support the two AFEs. The warm-end Cin::apse [63] connector assembly enables the connection between the flex circuits and the AFEs. The AFEs mate to backplane connectors mounted on a backplane support structure via card edge connectors. The backplane support structures are bolted to the VLPC cryostats, and thus the combination of cassettes and cryostat serve as the crates for the AFEs that are mounted on the cassettes. Note that this design (with connections on two orthogonal edges of the AFE) allows the AFEs to be removed for service without extracting the cassette from the cryostat.

The CFT requires 76,800 channels of VLPC readout, and the central and forward preshower detectors are instrumented with an additional 22,564 channels of VLPC readout. The ninety-nine cassettes necessary to provide this readout are mounted in two custom-designed cryostats located on the platform beneath the DØ detector. Better than 99.8% of the individual VLPC channels in these cassettes met or exceeded the desired performance specifications during cryogenic qualification tests performed prior to installation at DØ.

2.2.5 CFT readout electronics

The CFT and the central and forward preshower detectors (Section 4) share a VLPC-based readout and similar Level 1 and Level 2 trigger electronics, and therefore all three use the same front end electronics to process the signals from the VLPCs. The front end electronics are custom printed circuit boards

(the AFEs) approximately 14" tall (9U) and 18" long which are mounted on the VLPC cassettes inserted into cryostats as described in Section 2.2.4.

The AFE is a large and complex board that must perform a number of functions with competing requirements. It has charge-sensitive amplifiers to deal with the very small signals from the VLPCs. It is part of the Level 3 readout, part of the Level 1 and Level 2 triggers, and part of the slow control and monitoring system. It must also control the bias and temperature of the VLPCs. This functionality must be embedded in the AFE because it is the only piece of electronics that interfaces to the VLPCs. It must serve three different sub-detectors with different dynamic range requirements and comes in two types — left handed and right handed, depending on which side of the cassette body the board is mounted.

A number of features make it possible for the AFE to fulfill all of the above requirements. First, only a single printed circuit board was designed and laid out. The “handedness” of the board is determined by the way a few key components are installed. So, for example, there are mounting points and traces for the hard metric connectors on both ends of the board. When the connector is mounted on one end of the board, the board becomes left handed and mounts on the left side of the cassette. When the connector is mounted on the other end of the board, the AFE is right handed. The great majority of the components are mounted exactly the same way on both left and right handed AFEs. Second, the most noise sensitive parts of the board, including the front end amplifier chips, are mounted on separate multi-chip modules (MCMs), each with its own regulator for the power and separate ground planes. The amplifier chips are wire bonded directly to these much smaller (3.5" by 1.5") subassemblies. This allows the very fine pitch required for wire bonding to be confined to only the MCM substrate. Otherwise the AFE boards would not be manufacturable. This also allows the MCM subassemblies to be tested separately from the AFEs and to be removed as required for repair or replacement. There are eight MCMs on the AFE, each serving sixty-four VLPC channels, to match the construction of the cassette modules as described in Section 2.2.4. The MCMs are intended to be identical. To accommodate the different dynamic ranges of the CFT, the central preshower detector, and the forward preshower detector, capacitive charge division is used. By properly sizing the AC coupling capacitors on the AFE, the charge seen by the amplifiers is adjusted to accommodate the required dynamic range with the same electronics.

To further reduce the cost and simplify the system, the front end amplifier and digitizer is the same chip used for the SMT. This also facilitates commonality further downstream in the readout chain. The SVXIIe and its readout are described in Section 2.1.3. On the AFE, the SVXIIe provides for the integration of the charge signals from the VLPCs, a pipeline for storing the signals

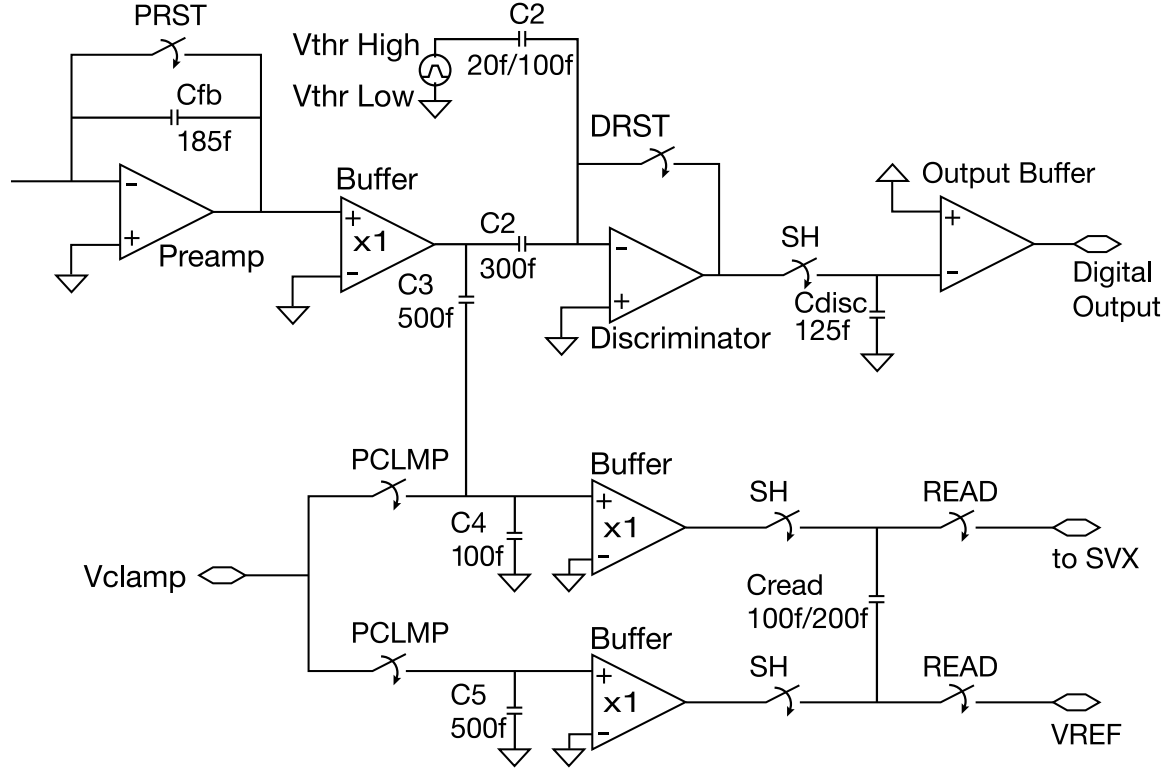


Fig. 19. Simplified schematic diagram of the SIFT chip. The bold text indicates clock signals supplied to the SIFT. Vclamp, Vthr, and VREF are externally-supplied DC voltage levels.

while the trigger is formed, and digitization of the signals and sparsification of the digitized data for readout. There are eight SVXIIe chips on every AFE, one on each MCM. The SVXIIe chips share a single 8-bit bus for readout. However, because the signals from the AFE are also needed for the trigger system, another chip, the SIFT chip, provides a trigger pickoff. A simplified schematic is shown in Figure 19. Each SIFT chip has eighteen channels so there are four SIFT chips on every MCM (eight channels are unused). The SIFT chips receive the signals from the VLPCs before the SVXIIe. For each channel, the SIFT has a preamplifier that integrates the incoming charge and switched capacitors that are used to split the amplified signal and send it along two paths: one to the SVXIIe, for later digitization and readout, the other to present the signal to a discriminator which fires if the charge exceeds a preset threshold (typically 10 to 15 fC).

There are five clock signals that control the operation of the SIFT, as can be seen from the schematic. One is the reset of the integrator, PRST, another is the reset for the discriminator, DRST. The third, SH (sample and hold) samples the discriminator output and the analog value of the integrator at the end of the cycle. READ transfers the analog charge to the SVXIIe. PCLMP resets the bias for the followers in the analog chain. All of these signals are derived from the main accelerator clock (53 MHz) by a combination of a DLL,

multi-tap analog delay, and programmable logic on the AFE board itself.

The digital signals from the discriminators serve as the inputs to the Central Track Trigger (CTT) and a bit is sent to the trigger system for every axial fiber every 132 ns (the time originally anticipated between beam crossings). This means that each AFE board is transmitting data at about 4.1 Gbits/s. The formatting and pipelining of this data is accomplished by sixteen CPLDs (complex programmable logic devices) in conjunction with four fast FIFO (first-in, first-out) memories. The trigger data is sent over four LVDS (low-voltage differential signaling) links. Each LVDS link is driven by Texas Instruments SN65LVDS95 transmitter chips [64]. Each chip has 21 bits of parallel-load, serial-out shift registers driven by a $7\times$ clock synthesizer and drives four LVDS lines (three line pairs are used to transmit 21 bits, the fourth carries the clock). With the data loaded into the transmitter chips with the 53 MHz clock, the AFE can send $21 \times 7 \times 4 = 588$ bits every 132 ns. One bit on each link is used to carry the 53 Mhz clock to help synchronize the data frames at the receiver end. This leaves $20 \times 7 \times 4 = 560$ bits for trigger information. The actual discriminator information is carried by 512 bits. The other bits are used to send control information to the CTT. These include indication of which crossings had beam present, when a Level 1 trigger accept was received from the trigger system, when the AFE is in a test mode or in reset, and so on. The control bits are duplicated on each of the four links and this redundancy is used to provide diagnostic information about the integrity of the links from the AFE to the CTT.

There are twenty CPLDs on each AFE. Besides the sixteen used to handle the discriminator data, one CPLD is used as a “virtual” SVXIIe chip to allow stored discriminator data to be read out on the same bus as the SVXIIe chips. Another CPLD is used to generate all the signals required to operate the SIFT chips and other timing signals for the board. The last two CPLDs work in conjunction with a microcontroller on the AFE to implement communications with the monitoring and slow control system via the 1553 bus. The microcontroller is a mixed signal controller PIC14000 from Microchip [65]. It has an integrated single-slope 16-bit ADC which is used to monitor analog values on the board such as VLPC bias, VLPC current, and VLPC temperature. The temperature of the VLPCs is measured via a calibrated carbon resistor mounted near the VLPC chips on every module. The value of the resistance must be measured to better than one part in a thousand and this drives the requirement for the ADC. The microcontroller implements closed-loop control for individual heater resistors on each VLPC module which allows the temperature at the VLPCs to be kept constant to 9.00 ± 0.05 K despite larger cryostat temperature fluctuations. The VLPC temperature monitoring and control is performed by the right-handed AFE boards. The VLPC bias voltage (which varies between 6 and 8 volts from module to module) must be controlled to approximately 50 mV to achieve optimal detector performance.




Fig. 20. A typical LED spectrum for a single VLPC for an axial CFT fiber. Every channel is fit automatically and the parameters of the fit are extracted and used for offline reconstruction. Typically, more than 97% of the axial channels are fit successfully. The solid histogram is the data; the smooth curve is the fit.

The extensive use of programmable logic on the AFE (twenty CPLDs and a microcontroller) greatly eased the task of designing, building, and operating it. It enabled the design of the AFE to be completed before all details of the interconnected detector systems were finalized. It allowed the fabrication of a single printed circuit board that is flexible enough to instrument three very different detectors. And it allows this very complicated board to operate with very little external control — once set points and operating parameters are downloaded, firmware on the board controls the system using on board DACs and ADCs. It also allows us to modify the firmware to optimize the operation of the electronics as running conditions change.

The performance of the electronics is dictated by the physics requirements of the detector. The small-diameter scintillating fibers used in the CFT coupled with the long waveguides necessary to direct the signals to the platform cause the signals generated by the VLPCs to be small (≈ 10 photoelectrons incident upon the VLPCs). To assure acceptable efficiency for triggers and tracking, the individual channel thresholds must therefore be set between 1.5 and 2 photoelectrons (pe). To maintain a low and stable threshold and to be able to distinguish individual photoelectron peaks during calibration, the analog signal must be digitized with a noise of less than 0.4 pe or about 2 fC. In fact, the front end electronics achieve or exceed all requirements. The mean pedestal width from fits to LED calibration spectra (Section 2.2.6) for all axial fibers is 0.24 pe or 1.6 fC and discriminator thresholds have similar noise and offsets of less than 2 fC so that it is possible to set discriminator thresholds below 10 fC for most channels. A sample spectrum from calibration data is shown in Figure 20. A summary of the important parameters for the AFE is given in Table 3.

Table 3
Summary of AFE board characteristics

AFE overview	
Number of readout channels per AFE	512
Number of MCMs per AFE	8
Total number of parts on AFE	2300
Threshold settings (analog V generated by onboard DACs)	4 per MCM
Pedestal settings (analog V generated by onboard DACs)	1 per MCM
Output bandwidth (trigger)	4200 Mbit/s
Output bandwidth (SVXIIe bus)	159 Mbit/s
AFE performance	
Noise (input referred) — analog readout	1.6 fC
Noise (input referred) — discriminator	1.2 fC
Threshold spread (rms)	2 fC
Pedestal spread (rms)	3 fC
AFE stability and reliability	
Mean time between failures (estimated)	35000 hrs/board
Calibration stability	Better than 1%

The readout of the SVXIIe chips mounted on the AFE boards is coordinated via signals from the Sequencers and Sequencer Controllers. The sparsified data is transferred from the Sequencers to VME Transition Modules (VTM) in the moveable counting house via fiber optics cables. The VTMs transfer the data to buffers in VRBs in VME crates. A VRBC mounted in the same crate controls the assignment of buffers. A single board computer mounted in that crate serves as the interface to the higher level readout. The readout of the CFT and the preshower detectors is buffered in four VME crates. Details of the higher level readout system are provided in Sections 2.1.3 and 10.

2.2.6 Calibration

The CFT and the preshower detectors are each equipped [66] with fast Nichia [67] blue-emitting LED pulsers that can be used to perform several functions. The pulsers certify the proper connectivity and yes-no operation of the fiber channels and attached readout chain, provide a means for monitoring the stability of the VLPC readout over time, and calibrate the response of individual

channels or, more specifically, provide a relative channel-to-channel energy calibration of the fiber-VLPC system. This monitoring and calibration is done during detector commissioning or between collider runs. Although the function of the LED system is the same for all three detectors, the method of light injection differs for each.

The limited space at the ends of the CFT cylinders complicates injection of light directly into the ends of the scintillating fibers. However, a method has been devised to introduce light through the cladding walls and into the active core of the fiber. The method takes advantage of the fact that the 3HF dye in the CFT fibers has a small absorption in the blue wavelength range (≈ 450 nm). Blue light incident on several layers of the scintillating fiber can be absorbed in the top layer, exciting the 3HF dye and producing scintillation photons detectable by the VLPC. To distribute the light over the cylindrical geometry of the fiber tracker, optical flat panels consisting of a single thin ribbon of 100 (typically) clear fibers, each approximately $500\text{ }\mu\text{m}$ in diameter, are degraded in certain local regions so that light is emitted through their cladding walls. The fibers are bundled together at one end of the ribbon, potted, and finished to allow coupling with an LED. The flat optical panels are then mounted around every fiber cylinder near each end of the cylinder and driven by pulsers. The use of two bands provides not only a margin of safety against failure of an individual optical element but also allows comparison of the light from the two ends of the tracker and monitoring of any changes in the quality of the scintillating fibers over time.

Although similar flat optical panels can be adapted to the preshower detectors, the CPS and FPS use different designs. For the preshower detectors, injection of LED light directly into the scintillator strips complicates the design appreciably, so instead, light is injected into the wavelength shifting (WLS) fibers. Since the WLS fibers are grouped in sets of sixteen channels and routed from the scintillator strip to connectors located at the perimeter of the detectors, one LED illuminates a group of sixteen WLS fibers. The mechanism, however, for delivering light into these sixteen fibers differs for CPS and FPS.

In the CPS, two LEDs are mounted inside a machined Lexan light block, directly before the outer perimeter of the detector, one on either side of the connector holding the WLS fibers. Ten LEDs mounted on one side of the connector are chained in series, allowing for an RG-178 cable to supply power to each LED string. The series of grouping ten LEDs continues across the perimeter of the CPS detector.

LEDs for the FPS system [68] are also mounted on the outer perimeter of the detector. However, due to the tight spatial constraints inside the FPS scintillating wedges, a 0.835-mm -diameter clear fiber transports the blue light to a group of sixteen WLS fibers, uniformly embedded within a hollow, cylindrical

barrel-cavity. The fiber is inserted directly into the central hole of the cavity, allowing for a relatively calibrated amount of light to be delivered over a full solid angle to the sixteen WLS fibers. The cavity holds the same group of WLS fibers that are routed from their respective scintillator strip to optical connectors coupling the clear fiber waveguides. The LEDs themselves are mounted on two PC boards, each providing the relevant electrical circuitry to drive the device.

PC boards located on the readout platform supply power to each detector's LED system. The entire system is computer-controlled and can be operated directly from the control room, allowing users to program the pulse-height and width as well as operate one LED in a detector at a time or the complete array. Distinct single photoelectron spectra have been observed and measurements made with the system have shown that connector-to-connector uniformity better than 20% can be achieved.

3 Solenoidal Magnet

The superconducting solenoidal magnet was designed [69] to optimize the momentum resolution, $\delta p_T/p_T$, and tracking pattern recognition within the constraints imposed by the Run I detector. The overall physical size of the magnet was determined by the available space within the central calorimeter vacuum vessel: 2.73 m in length and 1.42 m in diameter. We selected a central field of 2 T after considering the momentum resolution and tracking pattern recognition, the available space, and the thickness of the cryostat which depends on the thicknesses of the conductor and support cylinder. In addition, the magnet is required *i)* to operate safely and stably at either polarity, *ii)* to have a uniform field over as large a percentage of the volume as practical, *iii)* to be as thin as possible to make the tracking volume as large as possible, *iv)* to have an overall thickness of approximately $1X_0$ at $\eta = 0$ to optimize the performance of the central preshower detector mounted on the outside of the solenoid cryostat, and *v)* to quench safely without a protection resistor (although one is installed to shorten the recool time after an inadvertent fast dump). Services such as cryogenics, magnet high current buses, and vacuum pumpout and relief must reach the magnet from the control dewar through the narrow space (7.6 cm) between the central and end calorimeter vacuum vessels. The magnet system is controlled remotely, including cool down, energization, de-energization for field reversal, quench recovery and warmup, without access to the magnet cryostat, service chimney, or control dewar.

The major parameters of the solenoid design are listed in Table 4. A perspective view of the solenoid inside the central calorimeter with its chimney and control dewar is shown in Figure 21.

Table 4
Major parameters of the solenoid

Central field	2.0 T
Operating current	4749 A
Cryostat warm bore	1.067 m
Cryostat length	2.729 m
Integrated field homogeneity	$\pm 5.0 \times 10^{-3}$
Stored energy	5.3 MJ
Inductance	0.47 H
Cooling	Indirect, 2-phase forced flow helium
Cold mass	1460 kg
Conductor	18-strand Cu:NbTi, cabled
Conductor stabilizer	High purity aluminum
Thickness	0.87 X_0
Cooldown time	≤ 40 hours
Magnet charging time	15 minutes
Fast discharge time constant	11 seconds
Slow discharge time constant	310 seconds
Total operating heat load	15 W plus 0.8 g/s liquefaction
Operating helium mass flow	1.5 g/s

The solenoid, along with its cryostat, control dewar, and connecting service chimney, was manufactured by Toshiba Corp. [70] in Yokohama, Japan. The system was specified to operate safely and reliably over a twenty-year lifetime with up to 150 cool-down cycles, 2500 energization cycles, and 400 fast dumps.

3.1 Magnet construction

The solenoid is wound with two layers of superconductor to achieve the required linear current density for a 2 T central field. The support cylinder is located on the outside of the windings to support the radial Lorentz forces on the conductor, to provide axial rigidity to the finished coil, to provide a high thermal conductivity path to the helium coolant piping, and to ensure reliable quenchback for quench safety. To maximize the field uniformity inside the bore of the magnet, the current density in the windings is larger at the ends of the coil. This is accomplished by the use of narrower conductor at the ends




Fig. 21. Perspective view of the solenoid inside the central calorimeter. One end calorimeter and several muon chambers have been omitted for clarity. Also shown are the service chimney and control dewar.




Fig. 22. Cross sections of the two conductors used in construction of the solenoidal magnet.

of the coil. Cross sectional views of the conductors are shown in Figure 22. Both conductors are made with a superconducting Rutherford-type cable of multifilamentary Cu:NbTi strands stabilized with pure aluminum. The basic strand has a Cu:NbTi ratio of 1.34:1 and a diameter of 0.848 mm; eighteen strands are used in each conductor.

Both conductors are used in both layers. The middle section of each layer is wound with the wider conductor and the end sections with the narrower conductor. The transition point between the two conductors in the inner layer occurs at $z = \pm 0.941$ m, and at $z = \pm 0.644$ m in the outer layer. At the four locations where the conductor width changes, the two types of conductor are joined with a soldered lap joint one full turn in length. There are no other joints in the coil.

3.2 Magnet cryostat

The magnet cryostat consists of four major components: the vacuum vessel, the liquid-nitrogen-cooled radiation shield, the cold mass support system with liquid-nitrogen-cooled intercepts, and the helium cooling tube on the outer support cylinder of the superconducting coil.

The vacuum vessel consists of inner and outer coaxial aluminum shells with flat annular bulkheads welded to each end. The superconducting buses from the coil and the cryogen pipes from the outer support cylinder and the radiation shields leave the vacuum vessel through the service chimney nozzle welded in the bulkhead at the south end of the cryostat. The cryostat is designed for full internal vacuum and for an internal relieving pressure of 0.44 MPa (6.4 psig).

The magnet cold mass, the superconducting coil and outer support cylinder, weighs 1.46 metric tons. The cold mass support system consists of axial members which locate the coil axially and support it against axial thermal, decentering, and seismic forces, and nearly tangential members that locate the coil radially and provide radial support against thermal, gravitational, seismic, and decentering forces. The support members connect the outer support cylinder of the coil to the flat annular bulkheads of the vacuum vessel. All cold mass supports have thermal intercepts that operate near 85 K. The magnet cryostat is attached to the central calorimeter by support brackets that carry the weight of the cryostat and the tracking detectors which are attached to it, plus any residual magnetic decentering loads from the muon steel.

3.3 Control dewar and refrigeration system

The control dewar is the interface between the fixed cryogenic piping and the movable detector. It has bayonet connections for cryogenic lines and contains the vapor-cooled current leads and a liquid-helium reservoir, helium supply subcooler, helium flow control valve, and other valving and instrumentation. The current leads carry the full operating current of the magnet at the optimum cooling gas flow rate and are designed for safe operation without cooling gas flow for at least the full slow-discharge time constant of the magnet (approximately 310 s) in the event that cooling flow is lost while the magnet is energized.

Cryogenic and electrical services are carried from the control dewar to the magnet cryostat via the service chimney, which also serves as the relief line for the solenoid vacuum space and provides a path for pumpout of the magnet cryostat and control dewar vacuum spaces.

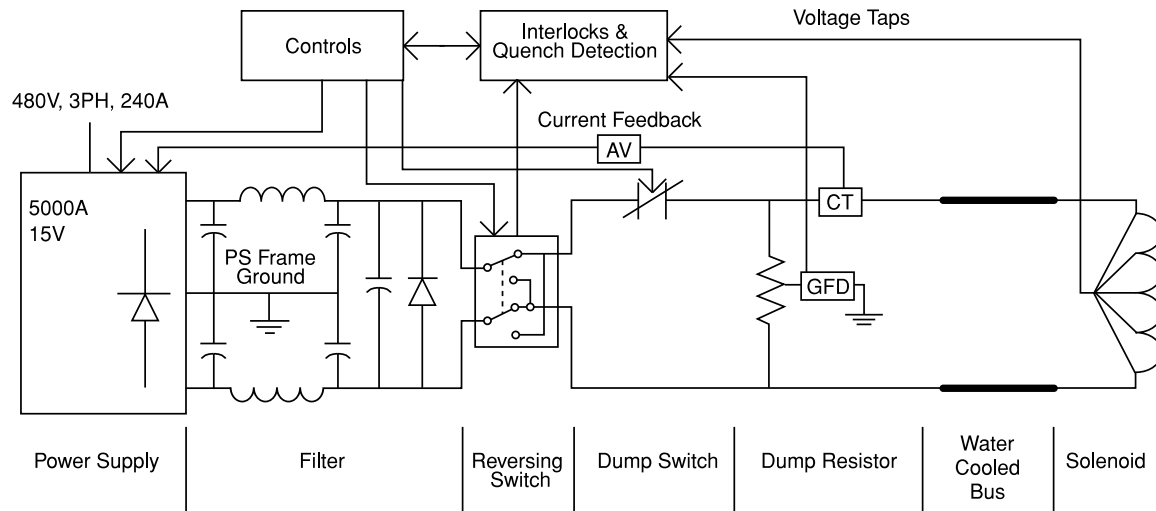


Fig. 23. Block diagram of the DC energization system for the solenoid.

A cryogenic refrigeration system supplies liquid nitrogen and liquid helium to the magnet and liquid helium to the VLPC cryostats for the CFT and central and forward preshower detector readouts (Section 2.2). The refrigeration requirements of the magnet and VLPC systems, as well as operation of the detector in the collision hall and adjacent assembly hall, were considered when the refrigeration system was designed.

A standard Fermilab satellite stand alone refrigerator (STAR) provides helium refrigeration. The capacity of the STAR is sufficient for non-simultaneous cooldown and simultaneous operation of both the solenoid and VLPC systems. Liquid helium is stored in a 3000 L dewar which supplies the magnet control dewar and the VLPC cryostats via separate transfer lines.

3.4 Magnet energization and control

A block diagram of the DC energization system is shown in Figure 23. The Power Energy Industries 150 power supply is a special Fermilab unit designed for superconducting loads. It is a twelve-phase-thyristor water-cooled rectifier unit with precision feedback current regulation. The power supply taps are set at 15 V/5000 A for efficient operation and reduced AC loading and DC ripple. The power supply regulates the current to within 0.01% using an external precision Holec 5000 A transducer installed downstream of the ripple filter and dump resistor.

Magnet polarity is reversed using a 5000 A DC mechanical motorized polarity reversing switch and a switch controller. Polarity reversal requires complete discharge of the solenoid followed by recharging at the opposite polarity. The controller confirms that the polarity reversal occurs at zero current load and

that the power supply is turned off. Polarity reversal takes about forty minutes and once the procedure is started, it must be completed before the original polarity can be restored.

The Texas Instruments TI565T Programmable Logic Controller (PLC) (originally installed to operate the liquid argon calorimeter refrigeration system, the muon chamber gas systems, and various building utility and safety systems) has been expanded to control and monitor the VLPC cryostats, solenoid, and helium refrigeration systems. The PLC is a stand-alone system with an internal battery-backed program memory which requires no host system to download the control programs in the event of a power failure. It consists of two independent TI565T processors operated in a “hot backup” configuration. One PLC is online and actively controlling the system; the second is in standby, running step-for-step with the online unit but with its input/output communication disabled. Each PLC runs continuous internal diagnostics for errors. When a fatal error is registered by the active PLC, it is taken offline and the standby PLC is put online without intervention or disruption to the system.

Dedicated control and monitoring of the magnet energization and protection system is done by a new Texas Instruments TI555 PLC. A quench protection monitor (QPM) hardwired chassis and interlock logic unit provide primary quench protection for the magnet. The QPM contains filtering, signal averaging, and delay circuitry for the voltage taps and thermometry which are used in the quench detection logic, vapor-cooled lead fault detection, or power supply failure logic. It contains preset limits for selected sensors which trigger fast or slow discharge of the magnet and it reserves the time ordering of the detected fault signals which trigger magnet discharge. The QPM defines a set of interlock conditions that must be satisfied before the magnet can be energized or the reversing switch operated.

The cryogenic controls and PLCs are powered by a 10 kW uninterruptable power supply that is backed up by an automatic diesel power generator. The total power consumption is approximately 6 kW.

3.5 *Magnetic field*

The magnetic field of the full magnet system is modeled using the TOSCA [71] program. The calculated field map was compared with the measured field in two locations: near the internal radius of the solenoid cryostat ($r \approx 54$ cm) at $z = 4$ and in the gap at the top of the central muon toroid steel. Within the solenoid, the measured field is 20.141 ± 0.005 kG; the calculated field at this location is 20.158 kG. The calculated magnetic field is scaled by 0.09% to




Fig. 24. $y - z$ view of the DØ magnetic field with both the toroid and solenoid magnets at full current. The field lines are projections onto the $y - z$ plane; the left and right line ends differ by up to 2.5 m in x . The numbers give the strength of the magnetic field at various locations in kG.

agree with the measurement. With full operating current in the toroid coils, here is a 4.5% difference between the calculated and measured field at the CF gap, requiring an adjustment in the calculated field for the CF toroid. A $y - z$ view of the magnetic field with both the toroid and solenoid magnets at full current is shown in Figure 24.

4 Preshower Detectors

The preshower detectors aid in electron identification and background rejection during both online triggering and offline reconstruction. They function as calorimeters as well as tracking detectors, enhancing the spatial matching between tracks and calorimeter showers [72]. The detectors can be used offline to correct the electromagnetic energy measurement of the central and end calorimeters for losses in the solenoid and upstream material, such as cables and supports. Their fast energy and position measurements allow preshower information to be included in the Level 1 trigger. The central preshower detector (CPS) covers the region $|\eta| < 1.3$ and is located between the solenoid and the central calorimeter. The two forward preshower detectors (FPS) [68] cover $1.5 < |\eta| < 2.5$ and are attached to the faces of the end calorimeters. The locations of the preshower detectors are shown in Figure 2.

4.1 Common properties of the preshower detectors

Both preshower detectors are made from triangular strips of scintillator, as shown in Figure 25. Since the triangles are interleaved, there is no dead space between strips and most tracks traverse more than one strip, allowing for strip-to-strip interpolations and hence improved position measurement. The strips are made by extruding polystyrene plastic doped with 1% p-terphenyl and 150 ppm diphenyl stilbene, with a light yield similar to that of commercial Bicron BC-404 scintillator. Each scintillator strip is machine-wrapped in aluminized mylar for optical isolation, and the end of each strip is painted white to enhance reflectivity. The packing density is different for the CPS and the FPS modules, resulting in different layer thicknesses and strip pitches. The average measured pitch differs somewhat from the design dimensions shown in Figure 25 and is 3.54 mm for the axial CPS layer and 3.65 mm for the interleaved FPS layers. After extrusion and wrapping, the triangular strips have a tendency to bow. To relieve stress in the plastic, making them easier to handle, the strips were slumped to the required shapes by heating to about 180° F for two hours followed by gradual cooling at room temperature.

Embedded at the center of each triangular strip is a wavelength-shifting fiber which collects and carries the light to the end of the detector. The non-readout ends of the WLS fibers are diamond-polished and silvered. At the readout end, fibers are grouped into bunches of sixteen and potted into connectors for transition to clear light-guide fibers. Light is transmitted via the clear fibers to VLPC pixels for readout. Both the WLS and clear fibers are 835- μ m Kuraray multicladd fibers.




Fig. 25. Cross section and layout geometry of the CPS and FPS scintillator strips. The circles show the location of the embedded wavelength-shifting fibers. Dimensions given are ideal, design dimensions.

The preshower detectors share common elements with the central fiber tracker, beginning with the waveguides and continuing through the entire readout electronics system. The last elements which are unique to the preshower detectors are the connections between the wavelength shifting fibers and the waveguides. For details on the waveguides, VLPC cassettes, and readout electronics, see Sections 2.2.3 – 2.2.5. The calibration systems for the preshower detectors are also similar to that of the CFT and are described in Section 2.2.6.

4.2 Central preshower detector

The CPS consists of three concentric cylindrical layers of triangular scintillator strips and is located in the nominal 5 cm gap between the solenoid and the central calorimeter. Between the solenoid and the CPS is a lead radiator 7/32" thick (approximately 1 radiation-length (X_0)) and 96" long, covering $|\eta| < 1.31$. The lead is covered by stainless steel skins 1/32" thick and 103" long. The solenoid itself is $0.9X_0$ thick, providing a total of about two radiation lengths of material for particles at normal incidence, increasing to about four radiation lengths at the largest angles allowed by fiducial constraints.

The three layers of scintillator are arranged in an axial- u - v geometry, with a u stereo angle of 23.774° and a v stereo angle of 24.016° . Each layer contains 1280 strips. The WLS fibers are split at $z = 0$ and read out from each end




Fig. 26. The CPS unwrapped in a plane. Each rectangle (trapezoid) is one octant module of the axial (stereo) layer. Note that the stereo octant edges are precisely aligned with the axial octant edges.

resulting in 2560 readout channels/layer. The channel count and geometry of the CPS matches that of the CFT for Level 1 triggering purposes. Each group of sixteen WLS fibers from the CPS axial layer corresponds to one of the eighty CFT sectors in ϕ . As with the CFT, the stereo layers are not used in the Level 1 trigger. However, unlike the CFT, the stereo layers of the CPS are used in the Level 2 trigger.

Each layer is made from eight octant modules. The modules consist of two $1/32$ " stainless steel skins with the scintillator strips sandwiched in between. The ends of the stereo modules align to the ends of alternating axial modules, as shown in Figure 26. The modules are attached directly to the solenoid by bolts at each corner. Eight $\frac{1}{4}$ " pins at each end of the solenoid provide additional registration. Connector blocks are spot welded between the stainless steel skins. The blocks provide structural integrity to the end region of the modules and mounting locations for the WLS connector and the cover for the light-guide connector.

4.3 Forward preshower detector

The two FPS detectors (north and south) are mounted on the spherical heads of the end calorimeter cryostats, occupying the region between the luminosity counters (Section 8) at the inner edge and the intercryostat detectors (Section 5.2) at the outer edge. Each detector is made from two layers, at different z , of a double layer of scintillator strips. A $2X_0$ -thick lead-stainless-steel absorber separates the two layers, as shown in Figure 27. The upstream layers are known as the minimum ionizing particle, or MIP, layers while the downstream layers behind the absorber are called the shower layers. All charged particles passing through the detector will register a minimum ionizing hit in the MIP layer, allowing measurement of the location (in η , ϕ , and z) of the track. Electrons will readily shower in the absorber, leading to a cluster of energy, typically on the order of three strips wide, in the shower layer that is spatially matched with the MIP-layer signal. Heavier charged particles are less likely to shower, producing a second MIP signal in the shower layer. Photons




Fig. 27. Complete ϕ -segment of a FPS module layer showing the overlapping $u-v$ MIP and shower layers, separated by an absorber.

will not generally interact in the MIP layer, but will produce a shower signal in the shower layer. The shower layers cover the region $1.5 < |\eta| < 2.5$ and the MIP layers and the absorber cover the region $1.65 < |\eta| < 2.5$. The outer region of the FPS, $1.5 < |\eta| < 1.65$, lies in the shadow of the solenoidal magnet coil, which provides up to $3X_0$ of material in front of the FPS. This amount of material readily induces showers that can be identified in the shower layers of the detector.

Each FPS detector has four measuring planes: MIP u and v and shower u and v . Each measuring plane is constructed from two layers, each containing eight 22.5° wedges (Figure 28) of active material separated by eight wedges of inactive material. The active material consists of two sublayers of nested scintillator strips with a crossing (“stereo”) of 22.5° with respect to one another. The two layers are separated in z ; the WLS fibers are brought to the periphery of the detector in the space between the layers where connection blocks facilitate the transition to clear waveguides which take the light to the VLPCs. Each MIP plane has 103 scintillator strips which are perpendicular to an edge at constant ϕ . Each shower layer has 144 strips, also perpendicular to a constant ϕ edge. Four special wedges in the vicinity of the solenoid cryogenics service pipes are notched to allow these to enter and have 142 strips per wedge in both the MIP and shower layers. The presence of these special wedges reduces the coverage to $1.8 < |\eta| < 2.5$ in this area. Successive FPS layers are offset to prevent projective cracks in ϕ . Figure 29 shows an $r-\phi$ view of two layers of the north FPS detector and Figure 30 describes the three distinct types of wedges and their dimensions within the successive FPS ϕ, z layers.




Fig. 28. Three-dimensional view of a FPS module mounted within its supporting frame.

The absorber layer is divided into forty-eight wedge-shaped segments for easier handling, each subtending 7.5° in ϕ and weighing approximately 5 lbs. Each segment consists of two lead absorber elements epoxied to each side of a $1/8$ "-thick stainless-steel plate. Each steel plate is approximately 0.5" longer radially than the lead plates, to allow connection to inner and outer support rings that are nested within the overall FPS layers. Similarly to the modules occupying the MIP and shower layers, the absorber segments are individually bent in three dimensions to conform to the spherical geometry of the end calorimeters. The ϕ edges of each segment are beveled to allow nesting of adjacent absorber segments with minimal non-projective gaps between them. The total thickness of a lead-stainless-steel-lead segment is 11 mm, or $2X_0$.




Fig. 29. $r - \phi$ view of the north FPS detector. For clarity, only layers 2 (shower) and 4 (MIP) are shown; layers 1 (shower) and 3 (MIP) are rotated by 22.5° in ϕ with respect to these layers. The south FPS detector is a mirror image of the north detector. The cutback near four o'clock accommodates the solenoid magnet cryogenic services.




Fig. 30. Relative orientation of two modules in successive FPS layers. The 12.7 mm overlap region is shown, along with the dimensions of the three types of modules.




Fig. 31. Isometric view of the central and two end calorimeters.

5 Calorimeters

The DØ calorimeter system consists of three uranium/liquid-argon calorimeters and the intercryostat detector.

5.1 Uranium/liquid-argon calorimeters

The calorimeters were designed to provide the energy measurement for electrons, photons, and jets in the absence of a central magnetic field (as was the case during Run I of the Tevatron), as well as assist in the identification of electrons, photons, jets and muons and establish the transverse energy balance in an event. The calorimeters themselves are unchanged from Run I and are described in detail in Ref. [43]. They are shown in Figure 31.

As shown in Figure 32, the central calorimeter (CC) covers $|\eta| \lesssim 1$ and the two end calorimeters, ECN (north) and ECS (south), extend coverage to $|\eta| \approx 4$. Each calorimeter contains an electromagnetic section closest to the interaction region followed by fine and coarse hadronic sections. The active medium for all of the calorimeters is liquid argon and each of the three calorimeters (CC, ECN, and ECS) is located within a cryostat that maintains the temperature at approximately 80 K. Different absorber plates are used in different




Fig. 32. Schematic view of a portion of the DØ calorimeters showing the transverse and longitudinal segmentation pattern. The shading pattern indicates cells for signal readout. The rays indicate the pseudorapidity intervals seen from the center of the detector.

locations. The electromagnetic sections (EM) use thin (3 or 4 mm in the CC and EC, respectively) plates, made from nearly pure depleted uranium. The fine hadronic sections are made from 6-mm-thick uranium-niobium (2%) alloy. The coarse hadronic modules contain relatively thick (46.5 mm) plates of either copper in the CC or stainless steel in the EC.

A typical calorimeter cell is shown in Figure 33. The electric field is established by grounding the metal absorber plates and connecting the resistive surfaces of the signal boards to positive high voltage (typically 2.0 kV). The electron drift time across the 2.3 mm gap is approximately 450 ns. Signal boards for all but the EM and small-angle hadronic modules in the EC are made from two 0.5 mm G-10 sheets. The surfaces of the sheets facing the liquid argon gap are coated with carbon-loaded epoxy [73] with a typical sheet resistivity of $40 \text{ M}\Omega/\square$; these surfaces serve as the high voltage electrodes for the gap. For one sheet, the other surface is bare G-10; the facing inner surface of the second sheet, originally copper-coated, is milled into the pattern necessary for segmented readout. Several such pads at approximately the same η and ϕ are ganged together in depth to form a readout cell. Calorimeter readout cells form pseudo-projective towers as shown in Figure 32, with each tower




Fig. 33. Schematic view of the liquid argon gap and signal board unit cell for the calorimeter.

subdivided in depth. We use the term “pseudo-projective” because the centers of cells of increasing shower depth lie on rays projecting from the center of the interaction region, but the cell boundaries are aligned perpendicular to the absorber plates.

There are four separate depth layers for the EM modules in the CC and EC. In the CC, the layers are approximately 2.0, 2.0, 6.8 and $9.8X_0$ thick. In the EC, they are approximately 0.3, 2.6, 7.9 and $9.3X_0$ thick. The material in the EC cryostat wall brings the total absorber for the first layer up to about $2X_0$ thick.

In the CC, the fine hadronic modules have three longitudinal gangings of approximately 1.3, 1.0, and $0.76\lambda_A$. The single coarse hadronic module has a thickness of about $3.2\lambda_A$. The two EC inner hadronic modules (Figure 31) are cylindrical, with inner and outer radii of 3.92 and 86.4 cm. The fine hadronic portion consists of four readout cells, each $17.6\lambda_A$ thick. The coarse hadronic portion has a single readout cell $4.1\lambda_A$ thick. Each of the EC middle hadronic modules has four fine hadronic readout cells of about $0.9\lambda_A$ each and a single coarse hadronic section of $4.4\lambda_A$. The outer hadronic modules of the ECs are made from stainless steel plates inclined at an angle of about 60° with respect to the beam axis (see Figure 31). The maximum thickness is $6.0\lambda_A$. Each layer of modules is offset from the others to prevent through-going azimuthal tracks.

The transverse sizes of the readout cells are comparable to the transverse sizes of showers: 1–2 cm for EM showers and about 10 cm for hadronic showers. Towers in both EM and hadronic modules are $\Delta\eta = 0.1$ and $\Delta\phi = 2\pi/64 \approx 0.1$. The third layer of the EM modules, located at EM shower maximum, is segmented twice as finely in both η and ϕ to allow more precise location of EM shower centroids. As can be seen in Figure 34, cell sizes increase in η and ϕ at larger η to avoid very small cells. Further details of the mechanical structure




Fig. 34. Layout of the EC electromagnetic readout cells for the four longitudinal EM layers. EM1 is closest to the interaction region. MH indicates the EC middle hadronic cells.

of the calorimeters can be found in Ref. [43]; these details are unchanged in Run II.

5.1.1 Calorimeter electronics

5.1.1.1 Overview There are 55,296 calorimeter channels to be read out; 47,032 correspond to channels connected to physical readout modules in the cryostats (the remaining channels are not connected). The readout is accomplished in three principal stages. In the first stage, the signals from the uranium/liquid-argon detector are transported to charge preamplifiers located on the cryostats via low impedance coaxial cable. In the second stage, the signals from the preamplifiers are transported on twisted-pair cable to the signal shaping and analog storage circuits (baseline subtractor boards or BLSs) located underneath the cryostats. The BLSs use switched capacitor arrays (SCAs) as analog storage devices to hold the signal for about $4 \mu\text{s}$ until the trigger is available, and provide baseline subtraction to remove any low frequency noise or pileup present in the signal. In addition, faster shaped analog sums of the signals are picked off to provide prompt inputs to the calorimeter trigger system for both the Level 1 and Level 2 trigger decisions. The precision signals from the BLSs are bused on an analog bus and driven by analog

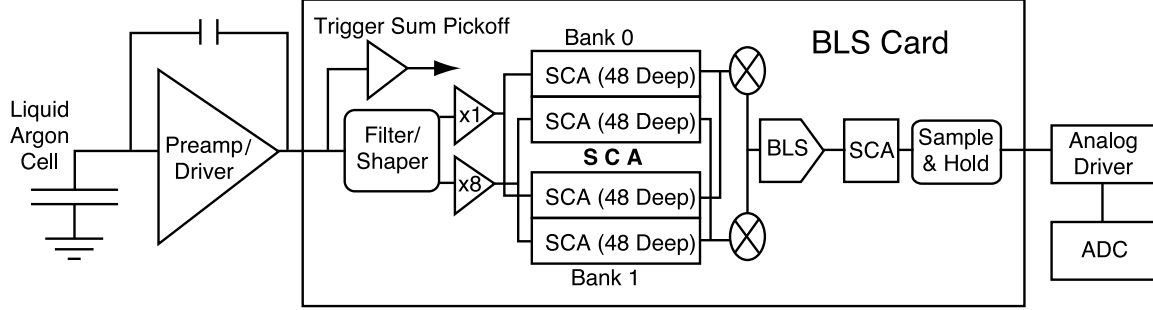


Fig. 35. Readout chain of the calorimeter in Run II indicating the three major components: preamplifiers, baseline subtractor and storage circuitry (BLS), and the ADCs.

drivers over 130 m of twisted-pair cable to analog to digital converters (ADCs) located in the moveable counting house, and then enter the data acquisition system for the Level 3 trigger decision and storage to tape. The preamplifiers and BLSs are completely new for Run II.

Figure 35 shows a simplified outline of the main components in the calorimeter readout chain.

5.1.1.2 Front-end electronics Signals from the calorimeter are transported on 30 Ω coaxial cable to a feedthrough port (the interface between the cold region and the warm region) on the cryostat and then on to the preamplifiers which are mounted as close as possible to the feedthroughs at the top of the cryostat. The cables from feedthrough port to preamplifier were replaced for Run II to provide better impedance matching to the preamplifier input ($\approx 30 \Omega$) and to equalize the lengths to provide better timing characteristics by minimizing the spread of the signal arrival. Electron drift time across the 2.3 mm liquid-argon gap remains approximately 450 ns at 2.0 kV for Run II, which provides a challenge for signal charge integration with beam crossings occurring every 396 ns. The calorimeter electronics was designed to maintain the good signal-to-noise ratio from Run I in the expected high instantaneous rate environment in the original Run II design, with a minimum bunch crossing timing of 132 ns.

The preamplifiers are individually packaged transimpedance (i.e. charge to voltage) hybrid amplifiers on ceramic substrates. Forty-eight individual preamplifier hybrids are mounted on a motherboard, with ninety-six motherboards housed in a single preamplifier box. Twelve such boxes, mounted on top of the calorimeter cryostats, make up the full calorimeter preamplifier system. Since access to these preamplifier boxes is very limited, the power supplies for the preamplifiers are redundant with each preamplifier box having two supplies, any one of which can be switched in to replace a failed supply. Low noise commercial switching power supplies are used to provide the necessary power

Table 5

Characteristics of the different preamplifier species used for the calorimeter and the ICD readout.

Preamplifier species	A	B	C	D
<Det. cap.> (nF)	0.26 – 0.56	1.1 – 1.5	1.8 – 2.6	3.4 – 4.6
Layer readout	EM1,2, HAD	HAD	HAD	HAD
Feedback cap. (pF)	5	5	5	5
RC (ns)	0	26	53	109
Total preamps	13376	2240	11008	8912
E	F	G	Ha-Hg	I
0.36 – 0.44	0.72 – 1.04	1.3 – 1.7	2 – 4	–
CC EM3	EC EM3,4	CC EM4, EC EM3,4	EC EM3,4	ICD
10	10	10	10	22
0	14	32	47-110	0
9920	7712	3232	896	384

density.

The preamplifier shaping networks are modified for cells of different capacitance to give approximately the same output waveform for all cells. Given the large range of detector capacitances at the input to the preamplifiers, there are fourteen species of preamplifier (plus one for readout of the intercryostat detector (ICD), Section 5.2) which provide similar output signal shapes into the BLS. This is important to maintain good timing for the peak-sampling circuit. The characteristics of the preamplifier species are shown in Table 5. The ICD feedback capacitor is 22 pF to lower the gain of the output signals to preserve their dynamic range.

Dual parallel input JFETs help maintain low noise performance. Two output driver stages provide the capability to drive a terminated 115 Ω twisted-pair cable. The power requirement of a single preamplifier is 280 mW. The noise performance can be evaluated from the width of the pedestal distribution, which is a function of the detector input capacitance.

The preamplifier motherboard is an eight-layer printed circuit board, with ground or power planes of solid copper separating planes containing signal traces to minimize noise pickup and cross-talk. The motherboard houses the precision resistors (0.1% 10 k Ω or 20 k Ω , depending on the preamplifier feedback capacitor) for the calibration voltage pulse. A single input line pulses six preamplifiers at once using a low capacitance trace. A discussion of the

calibration system is presented in Section 5.1.1.5.

All of the electronics is located in an area where there can be a residual magnetic field of a few hundred gauss (see Figure 24), and hence the switching supplies are magnetically shielded, and all other devices (including cooling fans) have been verified to function in the residual field. New heat exchange systems were installed in the existing preamplifier boxes to deal with the increased power dissipation.

5.1.1.3 Signal shaping and trigger The single-ended preamplifier signals are routed from the calorimeter over approximately 25 m of twisted-pair cable to the BLSs located below the cryostats where access is significantly easier. These boards perform two functions: the first is signal shaping, analog storage and baseline subtraction used for digitization; the second is a fast analog trigger sum signal for the Level 1 and Level 2 calorimeter triggers (Sections 9.1.2 and 9.3.1).

To minimize the effects of pile-up, only two-thirds of the charge collected by the preamplifier circuit, corresponding to the first ≈ 260 ns of signal collection from the gap, is used in the shaper circuit. The preamplifier output is an integral of the detector signal characterized by a rise time of about 450 ns and a recovery time of 15 μ s. Shaped signals are sampled every 132 ns, corresponding to seven RF buckets from the Tevatron. The shaper circuit produces an unipolar signal with a peak at about 320 ns and a return to zero after ≈ 1.2 μ s. The shaped signal is sampled at $t_0 = 320$ ns, close to the peak. To subtract the baseline, the signal three samples earlier ($t_0 - 396$ ns) is subtracted by the BLS circuitry.

The calorimeter Level 1 and Level 2 triggers are based on the energy measured in trigger towers of size 0.2×0.2 in $\Delta\eta \times \Delta\phi$, which is obtained by making appropriate sums (via resistor packs on the BLS boards) of the fast pickoffs at the shaper inputs. The resistor packs have been tuned based on a sampling weight optimization study which sought to minimize electron resolution first, and jet resolution second, for the channels included in the trigger (the coarse hadronic sections are not included in the trigger).

5.1.1.4 Digitization A BLS board contains forty-eight channels consisting of four towers with up to twelve preamplifier signals ganged to form a pseudo-projective calorimeter tower. There are signal shapers for each channel on the BLS motherboard and trigger pick-off and summation circuits tap the preamplifier signal prior to the shaper circuitry. The shaped preamplifier signals are fed to daughterboards, one per tower, each of which holds five SCA chips. The SCAs contain an array of 48 capacitors to pipeline the calorimeter

signals. The first and last buffers are not used in the readout to avoid edge effects in the chips. The SCA is not designed for simultaneous read/write operations, therefore two SCA banks are alternately employed for writing and reading the integrated charge. This scheme provides the $4.2\ \mu\text{s}$ buffering necessary prior to the arrival of the Level 1 trigger decision. There are also two gain paths ($\times 1$ and $\times 8$) to extend the ADC readout dynamic range so four of the SCAs are used to store the signals for the twelve channels on a daughtercard until the Level 1 trigger decision has been made. Once a positive Level 1 decision is received, baseline subtractor circuitry on each daughterboard decides channel-by-channel which gain path to use and subtracts the stored baseline (from the sample 396 ns earlier) from the peak signal. It then stores the result in the Level 2 SCA which buffers the data until after a Level 2 trigger decision. Once a Level 2 trigger accept is issued, the data are transferred from the Level 2 SCA to a sample-and-hold circuit on the daughterboard and an analog switch on the BLS motherboard buses the data on the BLS backplane to analog drivers which transfer the signal up to the ADC. The gain information is sent simultaneously on separate digital control cables.

The ADC successive approximation digitizers, reused from Run I, have only a 12-bit dynamic range, but the low and high gain path for each readout channel allow the 15-bit dynamic range. This matches the measured accuracy of the SCA. The readout system is approximately deadtimeless up to a Level 1 trigger accept rate of 10 kHz, assuming only one crossing per superbunch (group of twelve p or \bar{p} bunches, each bunch is separated by 396 ns; the Tevatron contains three evenly-spaced superbunches of each particle) is read out.

A master control board synchronizes twelve independent controllers (one for each readout quadrant of each calorimeter) in a shared VME crate. These controllers provide the timing and control signals that handle the SCA requirements and interface to the Level 1 and Level 2 trigger systems. Each control board houses three Altera FPGA (field programmable gate array) chips [74] (10K series, 208-pin packages). Three FIFO pipelines buffer up to about forty events awaiting readout at each trigger level. Timers on these event buffers ensure that “stale” data are appropriately flagged to the data acquisition system. Other readout errors are also flagged in the readout. The control boards permit the readout to be run in numerous diagnostic and calibration modes.

5.1.1.5 Calibration system The calorimeter calibration system consists of twelve identical units for the calorimeter and one slightly modified unit for the ICD. Each unit is composed of one pulser board and its power supply located in the BLS racks on the detector platform, and six active fanout boards housed inside the preamplifier boxes on the cryostats. The pulser boards are controlled via a serial bus to a VME IO register to set the amplitude and the timing of the calibration signal and to enable the channels to be pulsed.




Fig. 36. The mean current/DAC delivered by the pulser boards is 825 mA/DAC with a spread of $\pm 0.2\%$ for all 15 pulser boards produced (a) and $\pm 0.1\%$ within one board (b).

The pulser board delivers both a DC current corresponding to the chosen pulse height for each selected channel and a differential ECL command signal to the fanout boards. The pulse heights are set via an 18-bit DAC, with a maximum current of 100 mA delivered, and the timings through six programmable 8-bit delay lines with a nominal step size of 2 ns. The delays allows more careful adjustment of the relative timing of the calibrated charge injection and the readout of the preamplifiers. Sixteen switches located on each fanout board each generate a precise pulse on the reception of a command signal, converting the DC current to a calibration pulse, which is then distributed through the preamplifier box backplane to forty-eight calorimeter channels.

Both the pulser board and the active fanout have been shown to provide a pulser signal with a linearity at the per mil level and all the currents delivered are uniform within 0.2% between all boards and 0.1% within one board (Figure 36a). All the pulse shapes have been measured to estimate systematic effects on the signal amplitude, the timing, and the charge injected (Figure 36b).

5.1.1.6 Gain determination For all calorimeter channels, the gain calibration factors are determined from deviations of the slope ADC/DAC from its ideal value of 0.25 ADC counts per DAC unit. The dispersion of the coefficients for electromagnetic (hadronic) channels is about 5% (10%). Systematic shifts of the slope values can be observed for different preamplifier types.

Part of these differences is due to the injection of the calibration signal at the preamplifier input, outside of the cryostat. A fraction of the pulser signal travels down the signal cable and is reflected from the calorimeter cell, depending on the capacitance of the cell, and therefore has a different measured shape. The effects are largest for hadronic cells with a high capacitance, where differences in the timing can produce large tails in the distribution of the cal-




Fig. 37. Response to reflection measurement for three different channels. The vertical voltage scale is in arbitrary units. Cd is the detector cell capacitance.

ibration coefficients. Corrections for these effects have been derived for the calibration coefficients and the delay setting of the pulsed have been tuned to maximize the response in the electromagnetic layers.

Models of the electronics chain have been made to evaluate the differences between the electronics response to a calibration signal and a detector signal. To render these models realistic, all stable parameters of the signal path from the detector to the preamplifier input have been determined from reflection measurements. The reflected response (in arbitrary voltage units) to a step function is shown in Figure 37 for three different channel types. Quantitatively the values for the cable resistivity outside and inside the cryostat, the inductance of the feedthrough and the signal-strips as well as the capacitance of each cell have been determined and used in a simulation model. Correction factors to the calibration coefficients evaluate these effects at below the percent level for electromagnetic channels, when pulsed close to the signal maximum.

5.1.2 Liquid argon monitoring

The purity of the liquid argon is critical to the detector performance as electronegative contaminants (particularly oxygen) can combine with electrons traversing the gap and can severely impact the energy measurement. The liquid argon was recovered from Run I and, before refilling the cryostats for Run II, the purity was remeasured [75] with an external argon test cell equipped with ^{241}Am and ^{106}Ru radioactive sources. ^{241}Am is an alpha source (5 MeV α -particles, 432 yr half-life) that was used for purity measurements during Run I; ^{106}Ru is a new beta source (maximum 3.5 MeV β -particles,

≈ 1 yr half-life) with activity of 30–40 kBq. Calibration of the test cell was performed by injecting controlled amounts of O_2 into pure argon. The pollution of the liquid argon was measured to be less than 0.30 ± 0.12 ppm for all three calorimeter cryostats. A 5% signal loss is expected for a pollution level of 1 ppm.

Radioactive sources are used to monitor the contamination levels *in situ*. Each of the three cryostats is equipped with four ^{241}Am sources with an activity of $0.1 \mu\text{Ci}$ and four ^{106}Ru sources. Three of the beta sources in each cryostat now have very low levels of activity ($\lesssim 1$ Bq), fifteen years after the initial detector construction. One stronger source has an activity of about 4 Bq. The charge liberated in the liquid argon gap for the alpha sources is about 4 fC (about 25,000 electrons) and about half as much for the beta source. The expected trigger rate is about 500 Hz for the alpha source and 0.3 Hz for the strong beta source. The argon purity of the three calorimeters has been extremely stable over time.

A new readout board with preamplifier, dual operational-amplifier, and differential driver sends the amplified signals via a shielded twisted-pair cable to a differential receiver board to be digitized by a 12-bit ADC using a Xilinx FPGA (Spartan XCS40XL) [76]. Histograms can be accumulated very rapidly (few kHz) on the receiver board which are read out via CAN-bus to a PC (running National Instruments LabView [77]). This design is a modification of the liquid argon monitoring system for ATLAS and should be sensitive to less than 1 ppm of oxygen equivalent contamination.

The response of the signal is also a function of the liquid argon temperature so this is monitored as well. The temperature of the liquid argon varies by less than 0.1%.

5.2 Intercryostat detector

Since the calorimeter system is contained in three separate cryostats, it provides incomplete coverage in the pseudorapidity region $0.8 < |\eta| < 1.4$, as can be seen in Figure 32. In addition, there is substantial unsampled material in this region, degrading the energy resolution. To address this problem, additional layers of sampling have been added. Within the central and end cryostats are single-cell structures called massless gaps (calorimeter readout cells before the first layer of uranium). The ICD provides scintillator sampling that is attached to the exterior surfaces of the end cryostats. It covers the region $1.1 < |\eta| < 1.4$. Figure 2 shows the central region of the DØ detector, including the ICD.

The ICD, shown in Figure 38, is a series of scintillating tiles (Bicron BC-




Fig. 38. The arrangement of the ICD tiles on the endcap cryostats. The rectangles represent the drawers containing the ICD electronics. The beamline is perpendicular to the page.

400 [78]) connected to readout electronics devices via wavelength-shifting and clear plastic fibers (Bicron BCF-91A and BCF-98, respectively). The tiles cover an area of $\Delta\eta \times \Delta\phi = 0.3 \times 0.3$. Groups of three 0.9 mm wavelength-shifting fibers are embedded in grooves cut to a depth of 3.5 mm in the tiles. The wavelength-shifting fibers are mated to 1.0 mm clear optical fibers at the outer radius on the ICD tile enclosure. The clear fibers terminate at a Hamamatsu R647 photomultiplier tube [79] (PMT) connected to the back of the readout electronics. Each tile is subdivided into twelve readout tiles, isolated by grooves filled with epoxy mixed with highly reflective pigment. Because of the cryogenic services for the solenoid, there is a missing ICD tile in the south end of the detector; there is a total of 376 channels. The PMT then sends this signal through the electronics motherboard to the preamplifier. The preamplifier shapes the signal and sends it out of the ICD to BLS boards for processing.

Because of the solenoidal magnetic field in the DØ detector, the electronics of the ICD have to be located away from the tiles. Thus the container for the ICD electronics, the drawer system, is located in a low-field region to the side of the calorimeter. An ICD drawer houses PMTs, PMT electronics, and preamplifiers for six channels of readout. The readout signals are then multiplexed to interface with the BLS system common to the rest of the calorimeter. Internally, the ICD drawers have fibers impinging on the PMTs from the back end through the fiber backplane.

Clear fibers from the ICD tiles are mated to clear fibers contained in a light-tight enclosure called the “fiber backplane.” Between the fiber backplane and

the crate containing the drawers is a block of iron, with holes drilled to accept the PMTs. Fibers from the fiber backplane terminate in mounts in the iron. When a drawer is in the tray, six PMTs are aligned with six fiber bundles (two per channel plus one for a calibration signal) from the backplane. The PMTs are spring-loaded to ensure good contact with the fibers. The iron block, in combination with tubular magnetic shields, protects the PMTs from fringe magnetic fields in the region of the electronics drawers. O-rings at the base of the PMTs and the drawers themselves provide a dark environment for the PMTs.

The PMT high voltage dividers and preamplifiers are located on daughter cards that plug into the motherboard, which facilitates replacement. Input high voltage and output signal ports are located on the front of the drawer. Temperature monitors are placed between pairs of drawers to monitor heat dissipation from the preamplifiers.

The signal electronics are designed to be compatible with the DØ BLS/ADC system and electronics calibration. The ICD uses a modified version of the calorimeter preamplifiers, designed to stretch the PMT signal into a signal similar to that of the liquid argon readout. A modified form of the electronic pulser system is used to calibrate the preamplifier signals as well as to monitor performance of the electronics and verify that all channels are working. The responses of the PMTs are monitored using an LED calibration system, similar in design to the system used for the muon trigger scintillation counters (Section 6.4). An LED pulser is housed in each of the four fiber backplanes, and the level of the light signal is controlled externally.

Individual channel responses were measured on a test stand using cosmic rays. This test stand used the same combination of tiles, clear fibers, and electronics drawer that was mounted in the experiment after testing. Figure 39 show the distribution of cosmic ray peaks from the test stand. ICD sampling weights are taken from the full Monte Carlo simulation of the DØ detector, and are tuned using di-jet and photon-jet E_T balance.




Fig. 39. Distribution of all ICD channels' response to cosmic ray MIP signals.

6 Muon System

For muon triggering and measurement, the upgraded detector uses the original central muon system proportional drift tubes (PDTs) and toroidal magnets [43], central scintillation counters (some new and some installed during Run I), and a completely new forward muon system. The central muon system provides coverage for $|\eta| \lesssim 1.0$. The new forward muon system extends muon detection to $|\eta| \approx 2.0$, uses mini drift tubes (MDTs) rather than PDTs, and includes trigger scintillation counters and beam pipe shielding. The small angle muon system [43] of the original detector, including its associated magnets, has been removed.

During Run I, a set of scintillation counters, the cosmic cap [80], was installed on the top and upper sides of the outer layer of central muon PDTs. This coverage has been extended to the lower sides and bottom of the detector, to form the cosmic bottom. These trigger scintillation counters are fast enough to allow us to associate a muon in a PDT with the appropriate bunch crossing and to reduce the cosmic ray background. Additional scintillation counters, the $A\phi$ counters, have been installed on the PDTs mounted between the calorimeter and the toroidal magnet. The $A\phi$ counters provide a fast detector for triggering and identifying muons and for rejecting out-of-time background events.

The scintillation counters are used for triggering; the wire chambers are used for precise coordinate measurements as well as for triggering. Both types of detectors contribute to background rejection: the scintillator with timing information and the wire chambers with track segments.




Fig. 40. Exploded view of the muon wire chambers. **Note that this figure is being redrawn. Please do not comment on its quality, only on its content.**

New detectors and the modifications made to the original system are discussed in detail in the following sections; original components are described briefly. Exploded views of the muon system are shown in Figures 40 and 41.

6.1 Toroidal magnets

The toroidal magnets are described in detail in Ref. [43] and shown in Figures 1, 21, and 24. Having a stand-alone muon-system momentum measurement *i)* enables a low- p_T cutoff in the Level 1 muon trigger, *ii)* allows for cleaner matching with central detector tracks, *iii)* rejects π/K decays, and *iv)* improves the momentum resolution for high momentum muons.

The central toroid is a square annulus 109 cm thick whose inner surface is about 318 cm from the Tevatron beamline; it covers the region $|\eta| \lesssim 1$. To allow access to the inner parts of the detector, it was constructed in three sections. The center-bottom section is a 150-cm-wide beam, fixed to the detector platform, which provides a base for the calorimeters and central tracking detectors. Two C-shaped sections, which can be moved perpendicularly to the center beam, complete the central toroid. The magnet is wound using twenty coils of ten turns each. The two end toroids are located at $454 \leq |z| \leq 610$ cm. In the center of each end toroid is a 183 cm square hole centered on the beam-




Fig. 41. Exploded view of the muon scintillation detectors. Note that the backs of the scintillation detectors are shown for the south end. **Note that this figure is being redrawn. Please do not comment on its quality, only on its content.**

line; in x and y the magnets extend 426 cm from the beamline. The end toroid windings are eight coils of eight turns each.

During Run I, the central and end toroids' coils were operated in series at a current of 2500 A; the internal fields in the central magnet were approximately 1.9 T and those in the end magnets were approximately 2.0 T. During Run II, the magnets are again being operated in series, but at a current of 1500 A. The magnetic field is about 6% lower than that of Run I. Now, however, the primary measurement of the muon momentum is done using the new central tracking system and the reduced current provides a substantial cost savings. As in Run I, the polarity of the magnets during data collection is regularly reversed.

6.2 Central muon detector

The central muon system consists of a toroidal magnet, drift chambers, the cosmic cap and bottom scintillation counters, and the $A\phi$ scintillation coun-

ters.

6.2.1 Central muon drift chambers

The drift chambers are described in detail in Ref. [43]. The three layers of drift chambers are located inside (A layer) and outside (B and C layers) of the central toroidal magnet and cover $|\eta| \lesssim 1$. Approximately 55% of the central region is covered by three layers of PDTs; close to 90% is covered by at least two layers. The drift chambers are large, typically $2.8 \times 5.6 \text{ m}^2$, and made of rectangular extruded aluminum tubes. The PDTs outside of the magnet have three decks of drift cells; the A layer has four decks with the exception of the bottom A-layer PDTs which have three decks. The cells are 10.1 cm across; each chamber is 24 cells wide and typically contains 72 or 96 cells. Along with an anode wire at the center of each cell, vernier cathode pads are located above and below the wires to provide information on the hit position along the wire. The wires are ganged together in pairs within a deck and then read out by electronics located at one end of each chamber.

For each PDT hit, the following information is recorded: the electron drift time, the difference ΔT in the arrival time of the hit between a hit cell and the neighbor connected to it, and the charge deposition on the inner and outer vernier pads. Both ΔT and the charge deposition are used to determine the hit position along the wire. The drift distance resolution is $\sigma \approx 1 \text{ mm}$. The resolution of the ΔT measurement varies depending on whether the muon passes through the cell close to or far from the electronics. If the hit occurs far from the electronics, the resolution is approximately 10 cm. If it is close, the signal propagates two wire lengths and the dispersion in the signal causes the resolution to degrade to about 50 cm. Using charge division, the pad signal resolution is about 5 mm. However, only the A-layer pads are fully instrumented with electronics; about 10% of the B- and C-layer pads are instrumented (there are several reasons for this: *i*) for tracks traversing all three layers, the pad coordinate does not improve the pattern recognition or resolution significantly, *ii*) for tracks that only reach the A-layer, the additional information could help with track matching and background rejection, *iii*) the pad signals are used to monitor the gain to track aging in the PDTs — the instrumented B- and C-layer pads serve this purpose, and *iv*) fully instrumenting the B- and C-layer pads was too expensive).

To reduce the number of bunch crossings which occur during one maximum drift time interval, we are using a faster gas mixture than we used during Run I. The new mixture is 84% argon, 8% methane, and 8% CF_4 . The operating high voltage is 2.3 kV for the pads and 4.7 kV for the wires. The drift velocity is approximately $10 \text{ cm}/\mu\text{s}$, for a maximum drift time of about 500 ns. The contribution to the hit position uncertainty due to diffusion is about 0.4 mm,

worse than the 0.3 mm achieved using a slower gas during Run I. The poorer resolution is offset by the reduced occupancy and benefits to triggering due to decreasing the number of crossings in one drift time interval to two for 396 ns bunch spacing. The gas flow rate is 200 liters per minute, providing an exchange rate of three volumes per day. The gas is recirculated and filtered to remove contaminants.

Vapors from the copper-clad Glasteel (polyester and epoxy copolymer sheets with chopped glass fibers) [81] cathode pads are deposited on the wires in a sheath whose thickness is proportional to the accumulated integrated charge. As the coating thickens, the gain of the chamber drops and the chamber becomes inefficient. This aging during Run II has been significantly reduced by the increased gas flow, improved shielding from background particles, and removal of the Main Ring. The chamber wires could be cleaned if necessary, but no aging has been observed up to an integrated luminosity of 0.8 fb^{-1} . Since access to the A-layer chambers and the four central B-layer chambers directly under the detector is difficult, the cathode pads in these chambers have been replaced by copper-clad G10 and no aging is anticipated.

6.2.2 Cosmic cap and bottom counters

The cosmic cap and bottom counters are installed on the top, sides and bottom of the outer layer of the central muon PDTs. They provide a fast timing signal to associate a muon in a PDT with the appropriate bunch crossing and discriminate against the cosmic ray background.

The cosmic cap counters are described in detail in Ref. [80]. They are made from grooved 0.5" Bicron 404a scintillator with BCF 91a and Kuraray Y11 wave-shifting fibers glued into the grooves using Bicron 600 optical epoxy. There are 240 counters, 25" wide, and 81.5"–113" long. The counters are positioned with their width along z and length along ϕ . The grooves are 1.75 mm deep and 4 mm wide; they run along the length of the counter, from the end to just past the center. The grooves on each half of the counter are offset so they do not overlap at the center of the counter. They are spaced 8 mm apart so that half of the counter surface is covered with fibers. Each groove contains four fibers. The fibers are glued with five-minute epoxy at the ends of the counters and polished using a diamond cutter. To increase the light yield, a $1/32$ " anodized aluminum sheet is attached to the ends with aluminized mylar tape. The sides of the counters are milled. The fibers are gathered at the center of the counters, divided into two bundles with two fibers from each groove in one bundle and two in the other, and epoxied into two acrylic plastic cookies with holes in the center of each. The ends of the cookies are polished using the diamond cutter.

The scintillator is wrapped in a layer of Tyvek [82] with a hole for the fibers and cookies; around the hole is an aluminum lip. A layer of 1/8" thick Styrofoam is placed over the fibers on the counter top; aluminum sheets 0.020" thick cover the bottom and top surfaces. An outer frame of Unistrut provides support for the counter. A piece of black molded plastic fits over the outside of the aluminum lip and covers the phototubes, cookies and fibers.

The fibers are read out using two 1.5" 10-stage EMI [83] 9902KA phototubes mounted on the counters. The light yield varies depending on the distance from the phototube and the proximity to the counter edge. It is typically 30 photoelectrons per PMT for hits near the PMTs and 18 for hits near the distant corners.

The cosmic bottom counters complete the scintillator coverage of the central toroidal magnet. There are 132 counters, of two different designs. The forty-eight counters located on the outside of the center bottom B layer of PDTs (where there is no C layer) are nearly identical to the cosmic cap counters described above. Some minor improvements were made in the placement of the edge fibers to increase the light yield and the counter frames are made from 1/8" steel bent into U-shaped channels. The counters are suspended from the B-layer PDTs.

The sixty-eight counters located on the undersides of the remaining B and C layers of the PDTs are similar to the cosmic cap counters except that the bottom counters have fewer fibers and they are placed in vertical, rather than horizontal, grooves. The grooves are approximately 6 mm deep and 6–10 cm apart. This distribution of fibers results in the same light yield as the horizontal arrangement used for the cap counters. These counters use 1"-diameter MELZ [84] FEU-115M photomultiplier tubes. They are 12-stage PMTs with a 2 ns risetime and good quantum efficiency and uniformity. The PMTs are placed within 42-mm-diameter magnetic shields. The B-layer counters are suspended from the strong edges of the PDTs; the C-layer counters roll underneath the C-layer PDTs with one set of wheels in a track to maintain the counter position.

An important difference between the cosmic cap and cosmic bottom counters is that the bottom counters are positioned with their narrow dimension along ϕ and their long dimension along η . This orientation has better matching in ϕ with the central fiber tracker trigger. The counters' widths are approximately 4.5° in ϕ and they are as long or slightly longer than their respective PDTs are wide. Table 6 lists the location, number and size of the cosmic cap and bottom counters.

6.2.3 $A\phi$ scintillation counters

The $A\phi$ scintillation counters cover the A-layer PDTs, those between the calorimeter and the toroid. They provide a fast detector for triggering on and identifying muons and for rejecting out-of-time backscatter from the forward direction. In-time scintillation counter hits are matched with tracks in the CFT in the Level 1 trigger (Section 9) for high- p_T single muon and low- p_T dimuon triggers. The counters also provide the time stamp for low- p_T muons which do not penetrate the toroid and thus do not reach the cosmic cap or bottom counters.

An end view of the $A\phi$ counter layout is shown in Figure 42. The ϕ segmentation is approximately 4.5° which matches the central fiber tracker trigger (Section 9.1.3) segmentation. The longitudinal segmentation is 33.25" which provides the necessary time resolution and a match to the size of the PDTs; nine counters are required along the detector in the z direction. The nearly constant segmentation in ϕ is accomplished through the use of three sizes of counter: 14.46", 10.84", and 9.09" wide. The widest counters are located at the corners of the detector, the narrowest at the center of each side. There is a gap at the bottom of the detector where the calorimeter support is located. The counters overlap an average of about 3% in ϕ to reduce the possibility of muons passing through cracks. Along the length of the detector, the counters are butted end-to-end with a small gap between each.

The counters are made from Bicron 404A scintillator with G2 fiber embedded in vertical grooves [85] in a manner similar to the cosmic bottom counters described above. The grooves are about 1.75" apart and run along the length of each counter from the edge to the middle. Six fibers are spot glued into each groove and taper out of the groove at the middle of the counter. The fibers are routed to a MELZ FEU-115M photomultiplier tube which is secured to the counter case. The counter and fibers are wrapped in a layer of Tyvek, followed by a layer of black TEDLAR [49]. The counter case is an aluminum box with welded corners. It provides mechanical protection, support of the PMT and counter mounts, and protection against light leaks. The counters are mounted on aluminum cross-members attached to steel brackets which are fastened to the edges of the A-layer PDTs.

The $A\phi$ counters operate in a magnetic field of 200–350 G due to the residual magnetic field of the toroidal and solenoidal magnets. Magnetic shields made of 1.2-mm-thick μ -metal and 6-mm-thick soft iron with a diameter of 48 mm provide shielding from the fringe field (Section 6.3.2). The effect of the magnetic field on the PMT gain is less than 10% for any field direction.

The performance of an $A\phi$ prototype counter was studied using cosmic rays. The average muon signal corresponds to 50–60 photoelectrons. Comparison of




Fig. 42. End view of the layout of the $A\phi$ scintillation counters. The inset box shows an enlarged view of four counters. Azimuthal coverage is shown for seven of the counters. The bump on each counter represents the photomultiplier tube attached to the counter case.

the amplitudes of signals from the ends of the counters and from the center of the counter shows that the counters are uniform to $\pm 7\%$. The timing resolution of the counters is about 2 ns, due to photoelectron statistics, amplitude variation along the length of the counter, variation in the z -position of the vertex, variation in the z -position of the hit over the length of the counter, and variation of the time of collision. Different times-of-flight for particles at different polar angles are compensated for by varying cable lengths since the front-end electronics do not allow such timing adjustments.




Fig. 43. Cross-sectional view of a mini drift tube.

6.3 Forward muon system

The layout of the forward muon system is shown in Figure 1. It covers $1.0 \lesssim |\eta| \lesssim 2.0$ and consists of four major parts: the end toroidal magnets, three layers of MDTs for muon track reconstruction, three layers of scintillation counters for triggering on events with muons, and shielding around the beam pipe.

6.3.1 Mini drift tubes

Mini drift tubes were chosen for their short electron drift time (less than 132 ns), good coordinate resolution (less than 1 mm), radiation hardness, high segmentation, and low occupancy. The MDTs are arranged in three layers (A, B, and C, with A closest to the interaction region inside the toroidal magnet and C furthest away), each of which is divided into eight octants (Figure 40. A layer consists of three (layers B and C) or four (layer A) planes of tubes; each tube comprises eight $1 \times 1 \text{ cm}^2$ cells. The tubes are mounted along magnetic field lines (the field shape in the forward toroids is more “square” than “circular”). The entire MDT system contains 48,640 wires; the maximum tube length is 5830 mm in layer C. Since the flux of particles drops with increasing distance from the beam line, the occupancy of individual tubes is the same within a factor of two over an entire layer.

An MDT tube consists of eight cells, each with a $9.4 \times 9.4 \text{ mm}^2$ internal cross section and a $50 \text{ }\mu\text{m}$ W-Au anode wire in the center, see Figure 43. The tubes are made from commercially available aluminum extrusion combs (0.6 mm thick) with a stainless steel foil cover (0.15 mm thick) and are inserted into PVC sleeves. They are closed by endcaps that provide accurate positioning of the anode wires, wire tension, gas tightness, and electrical and gas connections. The anode wires are supported by spacers; the unsupported wire length never exceeds 1 m.

The MDT system uses a $\text{CF}_4\text{-CH}_4$ (90%-10%) gas mixture. It is non-flammable, fast, exhibits no radiation aging, and has a wide operational plateau. The maximum drift time for tracks which are perpendicular to the detector plane is 40 ns; for tracks inclined at 45° , the maximum drift time is 60 ns. Figure 44 shows the time-to-distance relationship for inclinations of 0° and 45° calcu-




Fig. 44. Time-to-distance relationship for a mini drift tube. The points labeled 0 degree and 45 degrees are calculated using GARFIELD. The crosses indicate measurements done at 0 degrees.

lated using GARFIELD [86] and for test beam data.

Negative high voltage is applied to the cathode (-3200 V); the anode wire is grounded at the amplifier. Each anode wire is connected to an amplifier and a discriminator located as close as possible to the detector. Each amplifier discriminator board (ADB) contains 32 channels and detects signals with a $2.0 \mu\text{A}$ threshold. Output logical differential signals from the ADB are sent to digitizing electronics which measure the signal arrival time with respect to the beam crossing with an accuracy of 18.8 ns. This large time bin (equal to the Tevatron RF bucket size) limits the coordinate resolution of the MDTs and was selected to reduce the cost of almost 50,000 TDCs. Hit information is sent to the trigger and data acquisition systems; the data acquisition system also receives drift times.

The efficiency of the MDTs is 100% in the active area of the cells for tracks which are perpendicular to the MDT plane. The overall plane efficiency is less, due to the wall thickness and PVC sleeves, and is approximately 95% for our geometry. The MDT efficiency is reduced near the wire-support spacers as well. Although each spacer is only 5 mm wide along the wire, the degradation in the electric field causes the efficiency to drop to approximately 20% for about 10 mm along the wire. An additional inefficiency of about 5% is caused by the tube endcaps and gaps between octants for mounting, gas connections, and HV and signal cables.

The momentum resolution of the forward muon spectrometer is limited by multiple scattering in the toroid and the coordinate resolution of the tracking detector. Although the MDT coordinate resolution measured in a test beam is about $350 \mu\text{m}$, the 18.8 ns time bin of the digitizing electronics leads to a resolution of about 0.7 mm per hit. The standalone momentum resolution of the forward muon system is approximately 20% for muon momentum below $40 \text{ GeV}/c$. The overall muon momentum resolution, including informa-




Fig. 45. Photograph of the C layer of muon trigger scintillation counters of the forward muon system.

tion from the silicon microstrip tracker and central fiber tracker, is defined by the central tracking system for muons with momentum up to approximately 100 GeV/c; the forward muon system improves the resolution for higher momentum muons and is particularly important for tracks with $1.6 \lesssim \eta \lesssim 2.0$, i.e. those which do not go through all layers of the CFT.

6.3.2 *Trigger scintillation counters*

The muon trigger scintillation counters are mounted inside (layer A) or outside (layers B and C) of the toroidal magnet (Figure 1). The C layer of scintillation counters is shown in Figure 45. Each layer is divided into octants containing about ninety-six counters. The ϕ segmentation is 4.5° and matches the CFT trigger sectors. The η segmentation is 0.12 (0.07) for the first nine inner (last three) rows of counters. The largest counters, outer counters in the C layer, are $60 \times 110 \text{ cm}^2$. The B and C layers have geometries similar to that of the A layer, but limited in places by the collision hall ceiling and floor.




Fig. 46. Design of a scintillation counter for the forward muon trigger system. Dimensions are in mm.

The counter design was optimized to provide good time resolution and amplitude uniformity for background rejection, high efficiency of muon detection, and reasonable cost for the production of nearly five thousand counters. A typical counter is shown in Figure 46. The counters are made of 0.5" Bicron 404A scintillator plate cut in a trapezoidal shape. Kumarin 30 [87] WLS bars are attached to two sides of the plate for light collection. The bars are 4.2 mm thick and 0.5" wide and bent at 45° to transmit light to a 1" phototube. The phototubes are fast green-extended phototubes, FEU-115M, from MELZ. They have a quantum efficiency of 15% at 500 nm and a gain of about 10^6 .

The sides of scintillator plates where the WLS bars are attached were left unpolished after milling. The unpolished sides give a larger number of photoelectrons and better uniformity than polished sides do and cost less [85]. The scintillator and WLS bars are wrapped in Tyvek for better light collection and black paper for light tightness. The wrapped counters are placed in 1-mm-thick aluminum containers with stainless steel transition pieces for connection of the phototube assembly.

The fringe magnetic field due to the solenoid and the toroidal magnet reaches 300 G in the region where the phototubes for layer A are located (Figure 24). To reduce the field in the area of the phototubes to approximately 1 G, they are placed in 48-mm-diameter magnetic shields made of 1.2-mm-thick mu-metal and 3- or 6-mm-thick soft iron. Since the field near the A layer is larger, the shields there use 6 mm of iron. Those for the B and C layers use 3 mm of soft iron. The effect of the magnetic field on the phototube signal is less than 10% for fields up to 350 G for the 6-mm-thick shields for any field orientation with respect to the phototube.

The performance of three counters of different sizes was studied in a 125 GeV/ c muon test beam. Figure 47 shows the dependence of the counter efficiency and time resolution as a function of high voltage for three counters: “large” 60×106 cm², “typical” 24×34 cm², and “small” 17×24 cm². A single discrimination threshold of 10 mV was used for these measurement. Time resolution of better than 1 ns and detection efficiency above 99.9% can be achieved at appropriate HV settings for all counter sizes. The non-uniformity of the counter response was measured by irradiating counters at different points using a ⁹⁰Sr source and checked by cosmic ray studies. For all counter sizes, the rms non-uniformity is less than 10%. Cosmic ray muons were used to determine the average number of photoelectrons detected in the counters. The largest counters have an average of 61 detected photoelectrons; the smallest counters give about three times as many photoelectrons.

Phototube signals are sent to 48-channel VME-based scintillator front end (SFE) cards. Input 1:1 transformers at the SFE card isolate the SFE from DC noise picked up between the phototubes and electronics modules which are located about 10 m apart. After amplification, the signals are sent to a 10-bit ADC and to a discriminator. Discriminated and gated signals are passed to the Level 1 trigger system and to the SFE TDC with a 1.03 ns time bin. After digitization, amplitude and time information is sent to the Level 2 trigger system and to the data acquisition system. The amplitude is measured in one out of sixteen channels per event for counter response monitoring.

The LED-based pulser system to monitor the forward muon trigger scintillation counters is described in Section 6.4.




Fig. 47. Time resolution and detection efficiency of muon trigger scintillation counters. See text for actual sizes.

6.3.3 Shielding

Three sources contribute to non-muon background in the central and forward muon systems: scattered proton and antiproton fragments that interact with the end of the calorimeter or with the beampipe produce background in the central and forward A layer; proton and antiproton fragments interacting with the Tevatron's low-beta quadrupole magnets produce hits in the B and C layers of the forward muon system; and beam halo interactions from the tunnel. Shielding installed in the accelerator tunnel during Run I [88] significantly reduced the background from beam halo. New shielding has been installed for Run II to reduce the background due to proton and antiproton remnants. Reduction in backgrounds along with the use of radiation-hard detectors guarantees the long-term, reliable operation of the muon system.

The shielding consists of layers of iron, polyethylene, and lead in a steel structure surrounding the beam pipe and low-beta quadrupole magnets. Iron is used as the hadronic and electromagnetic absorber due to its relatively short interaction (16.8 cm) and radiation (1.76 cm) lengths and low cost. Polyethylene is a good absorber of neutrons due to its high hydrogen content. Lead is used to absorb gamma rays.

The position of the shielding is shown in Figure 1. It extends from the end calorimeter cryostat, through the end toroid magnet, to the wall of the collision hall. It consists of three rectangular cross section pieces that are 84", 85" and 60" long, starting at the calorimeter and moving away from the detector center. The piece closest to the toroid has a 20" square hole at the center followed by 16" of iron, 6" of polyethylene, and 2" of lead. The two outer pieces are identical except that the hole is 25" square followed by 20" of iron. The most-forward section is split vertically down the center so that it can be moved

out of the way when the end toroid magnet is opened. Figure 45 shows the shielding in the “open” position.

Monte Carlo studies based on MARS code [89,90] show that the shielding provides a factor of 50–100 reduction in the energy deposition in muon detector elements. Reduction of backgrounds reduces the occupancy and detector aging, and provides almost background-free muon triggering and reconstruction. MARS Monte Carlo predictions of the number of hits in the muon detectors agree with the observed occupancies within 50%.

6.4 Muon scintillation counter monitoring

The central and forward muon scintillation counters are calibrated and monitored using an LED-based pulser system [91]. Given the large number of counters involved, it is difficult to use cosmic ray and beam muons to check the timing and do the PMT gain calibration quickly. The LED system allows us to find dead PMTs, isolate the behavior of the PMTs from the front-end electronics, adjust the relative timing between channels, set initial PMT gains, monitor PMT gains, and track timing changes.

A pulser triggers an LED driver board that drives four LEDs embedded in an acrylic block. The light is further mixed in two Lucite mixing blocks. The upstream side of each block is frosted by sand-blasting; the sides of the blocks are polished to maximize total internal reflection. The light is then distributed to an array of clear optical fibers embedded in a fourth acrylic block. These light distribution fibers are connected to fibers which are butted up to each PMT. To match the peak wavelengths and emission spectra of the wavelength-shifting fibers used in the counters, different LEDs are used for the central and forward counters. In the central region, blue-green NSPE510S LEDs from Nichia America Corp. are used; in the forward region, blue NSPB320BS LEDs, also from Nichia America Corp., have been installed. The clear fibers are Hewlett Packard HFBR-500 optical fibers.

The light intensity of the LEDs is monitored using a PIN diode (Hamamatsu S6775) mounted at the downstream end of the first mixing block; this works well because the mixing blocks evenly distribute the light to the fiber array. As long as the PIN diode is stable over time, variation in the gain of the PMTs can be measured independently of variation in the light output of the LEDs.

Figure 48 shows results from the forward muon system monitoring between July 2001 and January 2003. The ratio of the LED amplitudes is stable within 4% with a standard deviation of about 7%. The timing stability over this interval shows a peak stability within 0.2 ns with a standard deviation of about 0.43 ns. The typical time resolution of counters during data taking is

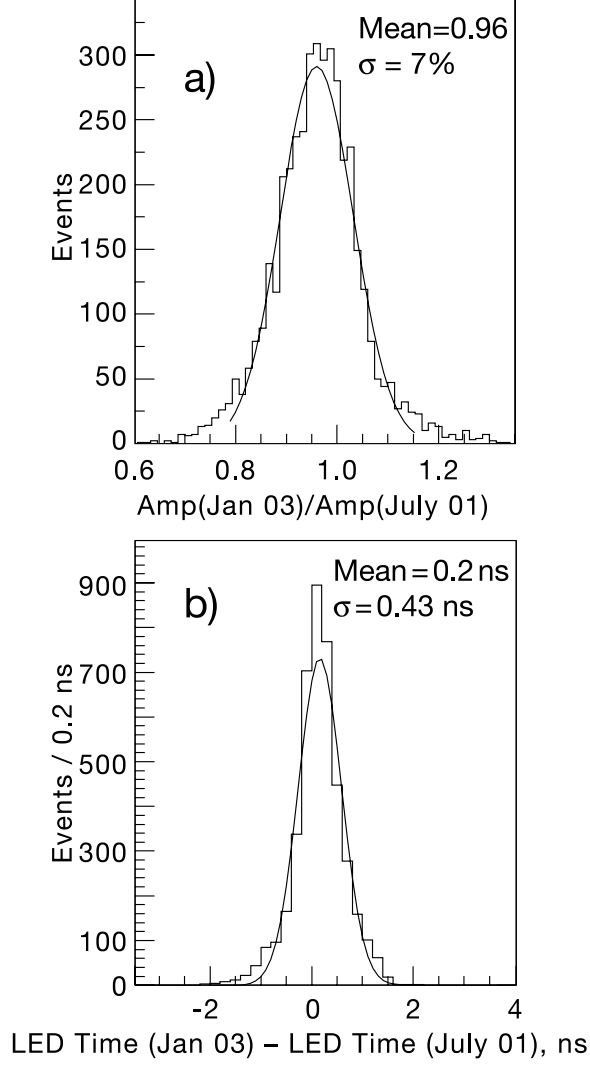


Fig. 48. Results from the forward muon system monitoring by the LED calibration system over one and one-half years; the smooth curves are Gaussian fits to the data. a) shows the LED amplitude ratio, b) shows the absolute difference in the LED timing.

2 ns, so stability is not a major factor in the overall time resolution of the system; photoelectron statistics, beam spot size, and bunch timing are all more important.

Table 6

Location, number and size of cosmic cap and bottom scintillation counters. All of the scintillator is 1/2" thick.

Location	Number	Width (inches)	Length (inches)	PMT
Cosmic cap top	80	25	113	EMI
Cosmic cap upper sides	80	25	108	EMI
Cosmic cap lower sides	80	25	81.5	EMI
Central side B layer	12	35	49	EMI
Central side B layer	4	35	67	EMI
Central bottom B layer	20	22 3/8	98 1/8	EMI
Central bottom B layer	20	15 3/4	98 1/8	EMI
Central bottom B layer lap	8	18 1/2	99 1/2	EMI
End bottom B layer	20	13 3/8	91 1/16	MELZ
End bottom B layer	12	19 1/4	91 1/16	MELZ
Central bottom C layer	20	22 1/16	68 1/16	MELZ
Central bottom C layer	16	29 3/10	68 1/16	MELZ




Fig. 49. FPD layout. Quadrupole castles are named P or A when placed on the p side or the \bar{p} side, respectively; U indicates up, D down, I in, O out.

7 Forward Proton Detector

The forward proton detector (FPD) [92] measures protons and antiprotons scattered at small angles (on the order of 1 mrad) that are missed by the main DØ detector. During Run I, such diffractive events were tagged using a rapidity gap (the absence of particles in a region of the detector), however a forward particle detector is necessary for access to the full kinematics of the scattered particle.

7.1 The detector

The FPD consists of a series of momentum spectrometers that make use of accelerator magnets in conjunction with position detectors along the beam line. The position detectors operate a few millimeters away from the beam and have to be moved away from the beamline during injection of protons into the accelerator. Special stainless steel containers, called Roman pots [93], house the position detectors, allowing them to function outside of the accelerator's ultra-high vacuum, but close to the beam. The scattered p or \bar{p} traverses a thin steel window at the entrance and exit of each pot. The pots are remotely controlled and can be moved close to the beam during stable conditions.

The Roman pots are housed in stainless steel chambers called castles. The FPD consists of eighteen Roman pots arranged in six castles. The castles are located at various distances from the DØ interaction point and in locations that do not interfere with the accelerator. The experimental arrangement of the FPD is shown in Figure 49. Four castles are located downstream of the low beta quadrupole magnets on each side of the collision point: two on the p side (P1 and P2) and two on the \bar{p} side (A1 and A2). Each of these quadrupole castles contains four Roman pots arranged to cover most of the area around the beam. Two castles (D1 and D2) are located on the outgoing \bar{p} side after the dipole magnet. Each of these dipole castles contains only one Roman pot.




Fig. 50. Detector fibers and the MAPMT.

7.2 The position detector

The construction of the position detectors is illustrated in Figure 50. Each detector is made of 0.8-mm-thick double-clad square scintillating fibers (Bicron BCF10 [78]) bundled in groups of four parallel fibers, forming a scintillating structure measuring $0.8 \text{ mm} \times 3.2 \text{ mm}$. One end of the detector element is aluminized (about a $3\text{-}\mu\text{m}$ -thick layer) to increase the light yield and the other end of each scintillating fiber is spliced to a double-clad clear fiber of square cross section (Bicron BCF98) with the same dimensions. The use of square fibers gives an increase of about 20% in light output compared to round fibers. The scattered p or \bar{p} goes through 3.2 mm of scintillating material. The four clear fibers take the light of one element to a single channel of the Hamamatsu H6568 16-channel multi-anode photomultiplier (MAPMT), yielding approximately ten photoelectrons.

As shown in Figure 50, each detector consists of six planes in three views (u , v and x) to minimize ghost hit problems and to reduce reconstruction ambiguities. Each view is made of two planes ($u-u'$, $v-v'$ and $x-x'$), the primed layers being offset by two-thirds of a fiber with respect to the unprimed layers. This arrangement yields a theoretical detector point resolution of $80 \mu\text{m}$. The u and v planes are oriented at $\pm 45^\circ$ with respect to the horizontal bottom of the detector, while the x plane is at 90° . There are twenty channels in each layer of the u and v planes and sixteen channels in each of the x layers. There are 112




Fig. 51. FPD castle.

channels (each with four fibers) per detector giving a total of 2016 channels in the eighteen Roman pots. Each detector needs seven MAPMTs and includes a trigger scintillator read out by a fast photomultiplier tube (Phillips XP2282 [94]). The FPD uses CFT electronics (Section 2.2.5) for the scintillating fiber detector read out, and luminosity monitor electronics (Section 8) for the trigger read out. The trigger manager designed for the muon system incorporates this information in making a Level 1 trigger decision on FPD tracks.

7.3 The castle

Figure 51 shows the FPD castle. It has four arms, each containing a Roman pot housing a detector. The castles are made of 316L stainless steel and, due to the ultra-high vacuum necessary in the interior of the castles, all parts were cleaned with demineralized water and alkaline detergent in an ultrasound bath and dried with a hot air jet before being TIG welded. The quadrupole castles (dipole castles are not in ultra-high vacuum) are baked at 150° C when the vacuum has been broken. A set of hot cathode and convection Pirani-style sensors monitors the vacuum in the chamber. Each castle has an associated ion pump to provide the ultra-high vacuum.

The castle sits on a stand that allows adjustment of its position in all directions over a range of 15 mm with an accuracy of 0.1 mm. The pot is connected to a driving system that makes it possible to move it perpendicularly to the beam. A 200- μ m-thick window separates the detector (inside the pot) from

the castle ultra-high vacuum. The system is operated by a step motor and a set of reduction gears allows pot motion with a precision of approximately 5 μm . A system of cylindrical and conical bearings allows adjustment of the pot alignment and a linear variable differential transducer (LVDT) monitors the pot position.

7.4 Acceptance and resolution

The FPD acceptance is maximized by minimizing the distance between the detectors and the beam axis. This distance is limited primarily by interaction with the beam halo which increases as the pots are inserted closer to the beam. FPD acceptance is determined as a function of t , the four-momentum transfer of the scattered proton or antiproton, and $\xi = 1 - x_p$ where x_p is the fractional longitudinal momentum of the scattered particle. For the dipole spectrometer, the acceptance is highest for $|t| \lesssim 2$, $0.04 \lesssim \xi \lesssim 0.08$ and extends to $|t| \lesssim 4.3$, $0.018 \lesssim \xi \lesssim 0.085$ (coverage is incomplete). The acceptance in the quadrupole spectrometers covers most of the region $0.6 \lesssim |t| \lesssim 4.5$, $\xi \lesssim 0.1$.

For elastic events both particles must be detected by diagonally opposite spectrometers with no activity detected in any other DØ subdetector. We used a sample of elastic events collected during special runs to measure the position resolution of the detector by using the u and v planes to measure the x coordinate and then compare it with the one given by the x plane. This process gives a detector resolution of 130 μm .

8 Luminosity Monitor

The primary purpose of the luminosity monitor (LM) is to make an accurate determination of the Tevatron luminosity at the DØ interaction region. This is accomplished by detecting inelastic $p\bar{p}$ collisions in a dedicated detector. The LM also serves to measure beam halo rates, to make a fast measurement of the z -coordinate of the interaction vertex, and to identify beam crossings with multiple $p\bar{p}$ interactions.

The luminosity \mathcal{L} is determined from the average number of inelastic collisions per beam crossing \bar{N}_{LM} measured by the LM: $\mathcal{L} = \frac{f\bar{N}_{LM}}{\sigma_{LM}}$ where f is the beam crossing frequency and σ_{LM} is the effective cross section for the LM that takes into account the acceptance and efficiency of the LM detector. Since \bar{N}_{LM} is typically greater than one, it is important to account for multiple $p\bar{p}$ collisions in a single beam crossing. This is done by counting the fraction of beam crossings with no collisions and using Poisson statistics to determine \bar{N}_{LM} .




Fig. 52. Schematic drawing showing the location of the LM detectors.

To accurately measure the luminosity, it is necessary to distinguish $p\bar{p}$ interactions from the beam halo backgrounds. We separate these processes by making precise time-of-flight measurements of particles traveling at small angles with respect to the beams using the LM detector (see Figure 52). We first assume that particles hitting the LM detector originate from a $p\bar{p}$ interaction and estimate the z coordinate of the interaction vertex z_v from the difference in time-of-flight: $z_v = \frac{c}{2}(t_- - t_+)$ where t_{\pm} is the time-of-flight measured for particles hitting the LM detectors placed at ± 140 cm. Beam-beam collisions are selected by requiring $|z_v| < 100$ cm, which encompasses nearly all $p\bar{p}$ collisions produced by the Tevatron ($\sigma_z \approx 30$ cm). Beam halo particles traveling in the $\pm \hat{z}$ direction will have $z_v \approx \mp 140$ cm, and are eliminated by the $|z_v| < 100$ cm requirement.

The LM detector consists of two arrays of twenty-four plastic scintillation counters with PMT readout. A schematic drawing of an array is shown in Figure 53. The arrays are located in front of the end calorimeters at $z = \pm 140$ cm, and occupy the region between the beam pipe and the forward preshower detector. The counters are 15 cm long and cover the pseudorapidity range $2.7 < |\eta| < 4.4$.

Scintillation light produced in the Bicron BC-408 scintillator is detected by Hamamatsu [79] R5505Q fine mesh PMTs (PMTs). Due to space constraints and the characteristics of the PMTs, they are mounted on the faces of the scintillators with the axes of the PMTs parallel to the z axis. They have no magnetic shielding, and their gain is reduced by a factor of about 30 when the solenoidal magnet is turned on due to the approximately 1 T magnetic field in this region [95]. The time-of-flight resolution for the counters is about 200 ps, with the dominant contribution to the resolution being the variation in light path length for particles striking different locations on the scintillator.

Radiation damage is a concern for detectors located this close to the beams. Much of the radiation dose seen by these detectors comes from the $p\bar{p}$ collision products, and is thus unavoidable. The PMTs are exposed to a radiation flux of about 25 kRad/fb^{-1} , which is sufficient to cause darkening of the borosilicate glass window typically used for PMTs. The R5505Q PMTs have fused silica (quartz) windows that are largely immune to radiation damage [79]. The radiation flux increases rapidly with decreasing radius, reaching a level of ap-




Fig. 53. Schematic drawing showing the geometry of the LM counters and the location of the PMTs (solid dots).

proximately 300 kRad/fb^{-1} at the innermost scintillator edge. Based on the radiation damage study in Ref. [96], modest ($\approx 10\%$) light loss is expected for the innermost scintillator edge after 3 fb^{-1} .

The scintillation counters are enclosed in light-tight enclosures, with each enclosure holding twelve counters. Preamplifiers inside the enclosures amplify the PMT signals by a factor of five. The fused silica PMT windows are much more permeable to helium gas than borosilicate glass [97]. To avoid damage from the widely fluctuating helium concentration present in the DØ collision hall, the enclosures are purged with dry nitrogen.

For accurate timing of the PMT signals, low-loss cables [98] are used to bring the signals from the detector to the digitization and readout electronics. The signals are equalized in time and split into two paths. On one path, analog sums are formed from the PMT signals for each of the two arrays, which are then timed using a digital TDC to identify $p\bar{p}$ collisions [99]. On the other path, two types of custom VME boards provide the required signal processing. Six LM-TDC boards are used to digitize the time and charge for each PMT and apply charge-slewing corrections to generate fully calibrated time-of-flight measurements. A single LM-VTX board utilizes the measurements made on the LM-TDC boards to calculate the average time for each counter array and the z -coordinate of the interaction vertex, identify beam crossings likely to have multiple $p\bar{p}$ interactions from the spread in time-of-flight measurements, and provide interaction and multiple-interaction terms to the Level 1 trigger.

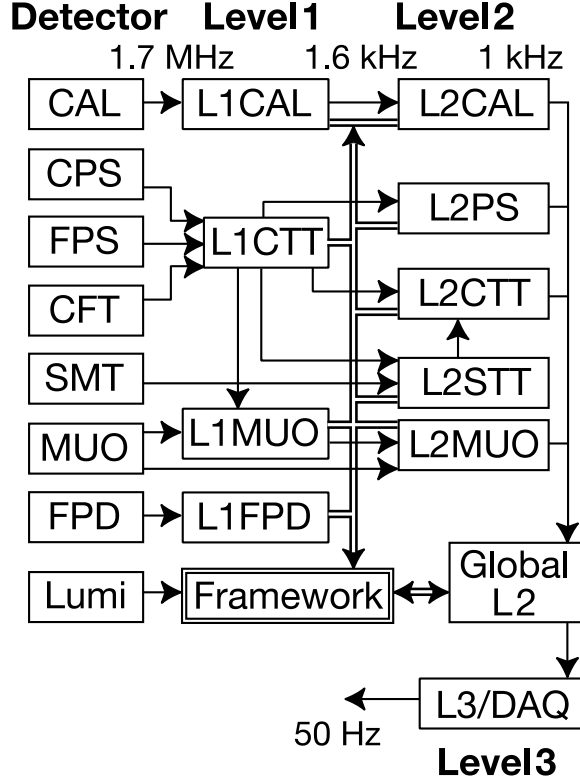


Fig. 54. Block diagram of the DØ trigger system. The arrows show the flow of trigger-related data.

9 Triggering

With the increased luminosity and higher interaction rate delivered by the upgraded Tevatron, a significantly enhanced trigger is necessary to select the interesting physics events to be recorded. Three distinct levels form this new trigger system with each succeeding level examining fewer events but in greater detail and with more complexity. The first stage (Level 1 or L1) comprises a collection of hardware trigger elements that provide a trigger accept rate of 1.6 kHz. In the second stage (Level 2 or L2), hardware engines and embedded microprocessors associated with specific subdetectors process information used by a global processor in determining correlations between different detectors. L2 decisions are made within $100 \mu\text{s}$ with an accept rate of about 1 kHz. Candidates from L1/L2 are passed to a farm of Level 3 (L3) microprocessors; sophisticated algorithms reduce the rate to about 50 Hz and these events are recorded for offline reconstruction. A block diagram of elements of the DØ trigger system is shown in Figure 54.

The trigger system is closely integrated with the read out of data, as illustrated in Figure 55. Each event that satisfies the successive L1 and L2 triggers is fully digitized, and all of the data blocks for the event are transferred to a single commodity processor in the L3 farm. The overall coordination and control

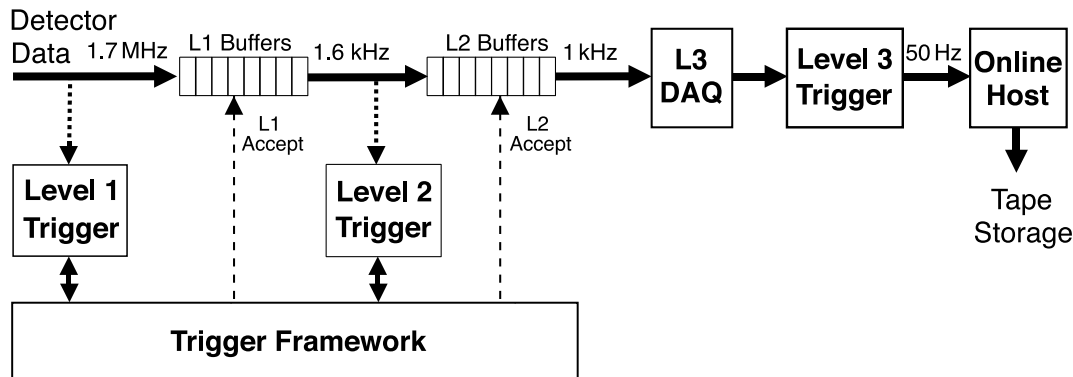


Fig. 55. Overview of the DØ trigger and data acquisition systems.

of DØ triggering is handled by the COOR package running on the online host. COOR interacts directly with the trigger framework (for L1 and L2 triggers) and with the DAQ supervising systems (for the L3 triggers). The data acquisition system responsible for the data flow of the fully digitized event into L3 is described in Section 10.

9.1 The Level 1 trigger

L1 is implemented in specialized hardware and examines every event for interesting features. The calorimeter trigger (L1Cal) looks for energy deposition patterns exceeding programmed limits on transverse energy deposits; the central track trigger (L1CTT) and the muon system trigger (L1Muon) compare tracks, separately and together, to see if they exceed preset thresholds in transverse momentum. The L1 forward proton detector trigger (L1FPD) is used to select diffractively produced events by triggering on protons or antiprotons scattered at very small angles.

All of the L1 triggers are pipelined and thus make minimal contributions to the deadtime and, in order to participate in the trigger decision, they must arrive at the trigger framework (Section 9.1.1) in $3.5 \mu\text{s}$ or less. The rate of L1 trigger accepts is limited by the maximum readout rates of the participating subsystems and by a desire to minimize deadtime associated with these readouts.

9.1.1 Trigger framework

The L1 trigger framework (TFW) gathers digital information from each of the specific L1 trigger devices and chooses whether a particular event is to be kept for further examination. In addition, it coordinates various vetoes that

control program for the experiment, via text-based commands interpreted in the TCC that sets the low level hardware registers.

- Any trigger can be prescaled under program control. The L1 TFW manages the prescale ratios.
- At the level of data acquisition, the detector is divided into subsections or groups of subsections served by a single serial command link (SCL) called geographic sections. Typical geographic sections are detector front-ends but other systems such as the L1 TFW readout are also geographic sections. The L1 TFW supports up to 128 geographic sections and any specific trigger can be programmed to request digitization and readout of any subset of the 128 geographic sections. This permits a partitioning of the data acquisition system into several separate systems, which is particularly useful during setup and installation. Each geographic section receives an L1Accept signal from the L1 TFW over the SCL, followed by an L2Accept or L2Reject from the L2 global processor.
- The L1 TFW can be programmed so that only particular bunch crossings contribute to a given trigger. Beam crossing numbers and L1Accept numbers provide a unique identifier for events.
- The L1 TFW provides all of the scalers needed to count the triggers as well as scalers to monitor the “live” time of combinations of these triggers assembled into exposure groups.
- The TFW offers an extensive menu of lower level and specialized commands for diagnosis of the trigger hardware by specialists.
- TFW scalars and registers can be accessed through the TCC to perform trigger system programming, diagnostics, and monitoring.
- During Run I, elements of the L1 trigger could only be ANDed to give a specific trigger. For Run II, the capability to allow both ANDs and ORs of trigger terms has been added to the TFW, using resident firmware. The resulting combinations are called pseudo-terms, since they are not formed in hardware. The pseudo-terms have the ability either to sharpen trigger turn-on curves, or to reduce combinatorial backgrounds and thus reduce the rate of low- p_T triggers not otherwise affordable.

9.1.2 Level 1 calorimeter trigger

The underlying architecture of the L1 calorimeter trigger for Run II is the same as it was for Run I [100]. There are, however, a number of improvements in the way the data are received by the trigger and the way the results are extracted.

9.1.2.1 Summary of operational principles The trigger inputs consist of electromagnetic (EM) and hadronic (H) trigger tower energies made up from sums in depth and transverse coordinates ($\Delta\eta \times \Delta\phi = 0.2 \times 0.2$) of fast analog

pickoffs from the BLS circuits. There are 1280 EM towers and 1280 H towers: forty slices in η , covering the region $|\eta| < 4$, and thirty-two slices in ϕ , covering the full 2π of the azimuth. The tower energies are converted to E_T on input, and have pedestals subtracted and energy scales adjusted if necessary. The variables used in trigger calculations are actually the EM transverse energies and the total transverse energies (EM+H) formed by adding the corresponding EM and H towers. The triggers available for use in the experiment include:

- Global variables
 - ΣE_T — with four thresholds
 - \cancel{E}_T — with four thresholds.
- Local variables

Each EM tower and each EM+H tower is compared to four programmable E_T threshold sets. A given threshold set could contain different values of E_T for each tower although in practice they are all set to the same value. A bit is set if a tower exceeds its reference value. (A feature vetoing the EM tower if its corresponding H tower exceeds a programmed reference value is available but has never been used.) The number of EM towers and EM+H towers exceeding their thresholds in a given reference set is counted and trigger bits are generated if this count exceeds any of several programmed count limits. We chose to use two such count limits for EM towers and four for EM+H towers.
- Large tiles

Since the trigger towers are small relative to the size of jets, we use some of the partial E_T sums needed for the \cancel{E}_T calculation to trigger on jets. Conveniently available are sums covering 4×8 trigger towers in $\eta \times \phi$, corresponding to an area of 1.28 in this space. For generating trigger bits, four reference sets are available with two count thresholds for each set but only for EM+H and not for EM-only large tiles.

9.1.2.2 Modifications for Run II A number of changes were made to the calorimeter trigger in anticipation of the higher crossing rates and luminosities expected during Run II.

- Signal receiver electronics

In Run I, the conversion of the BLS energy signals to properly-scaled E_T values was done on the calorimeter trigger front end (CTFE) cards by a network of hard-wired precision resistors, making gain adjustments difficult. For Run II, we bypassed this problem by building new front-end receiver cards capable of gain adjustment up to a factor of two for each channel under computer control. This change also allowed us to add circuitry to decrease the trigger signal rise times to be compatible with the 132 ns accelerator bunch spacing originally expected for Run II. Each card processes the signals from two trigger towers and two such cards were needed to upgrade each

- Run I CTFE card.
- Readout electronics

To speed up the readout of L1 calorimeter trigger data and therefore improve the L1 trigger accept rate, the old readout system of two eight-bit buses was replaced with ten sixteen-bit buses. With this improvement, the L1Cal accept rate can exceed 10 kHz.
- Trigger coverage

Trigger tower energy data are read out for the full calorimeter. Because of signal to noise considerations, only the trigger towers for $|\eta| < 3.2$ are used for triggering. In addition to the calorimeter trigger towers, ICD towers are embedded in the data as well.

9.1.3 Level 1 central track trigger

The L1CTT [101] reconstructs the trajectories of charged particles using fast discriminator data provided by the three scintillator-based detectors: the central fiber tracker (Section 2.2) and the central and forward preshower detectors (Section 4). Data processed by the CTT are used to make L1 trigger decisions, otherwise known as trigger terms. While the CTT is optimized for making fast L1 trigger decisions, the electronics also store more-detailed event data (e.g. sorted lists of tracks and preshower clusters) for later L2/L3 readout, or for use as seeds in other DØ trigger systems.

9.1.3.1 Overall system design The input to the CTT system consists of the discriminator bits generated on the AFE boards every 132 ns. These discriminator bits are sent from the AFE boards over point-to-point low voltage differential signal (LVDS) links to chains of digital front-end (DFE) boards. All of the DFE boards in the CTT system are built using a common 6 U \times 320 mm motherboard that supports one or two daughterboards. Currently there are two different daughterboard layouts with different sizes and numbers of FPGAs. A transition board allows a DFE motherboard to drive LVDS, fiber optic, and coaxial copper links. CTT-specific protocols define the data format for communication between all DFE boards and consumers. This hardware modularity, when coupled with the flexibility of FPGAs, enables application-specific functionality to be located in firmware, minimizing the number of unique DFE boards in the system.

The CTT system comprises three subsystems: CFT/CPS axial, CPS stereo, and FPS. Of these, the CFT/CPS axial and FPS subsystems make L1 trigger decisions and send trigger terms to the trigger framework. All three subsystems participate in L2/L3 readout by sending track and cluster lists to various preprocessor and readout crates with a deadtime not exceeding 5%. In the following sections we focus primarily on the CFT/CPS axial subsystem.

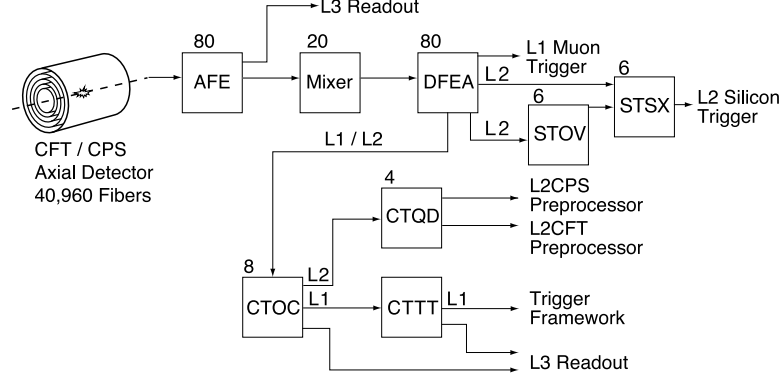


Fig. 57. The CFT/CPS axial components of the CTT system.

9.1.3.2 CTT subsystems

- CFT/CPS axial

The CFT/CPS axial subsystem (Figure 57) is designed to provide triggers for charged particles with $p_T > 1.5 \text{ GeV}/c$ at the highest possible efficiency. In addition to finding tracks, it must also find CPS clusters, match tracks to clusters, and report the overall occupancy of the CFT axial layers. Significant resources are allocated for triggering on isolated tracks. The CFT/CPS axial system also supplies the L1Muon and L2STT systems with lists of seed tracks, and sends track and cluster information to the L2CTT and L2PS preprocessors (Figure 54).

For mechanical reasons, the CFT and CPS axial fibers are grouped by cylinder layer before routing to the AFE boards. However, the track finder algorithms require the fiber information to be arranged in 4.5° sectors in the transverse plane (Figure 58). Furthermore, each track-finder daughterboard must receive information from each of its neighboring sectors to find tracks which cross sector boundaries. A set of twenty mixer boards [102] handles this data reorganization and duplication. After data duplication, the mixer output data rate is a constant 475 Gbits/s. Total latency through the mixer system is 200 ns.

From the mixer system, the properly organized discriminator bits are sent to the first tier of DFE boards. This first tier consists of forty motherboards, with each motherboard having two singlewide daughterboards (DFEAs). Each DFEA daughterboard unpacks the CFT data and compares the fiber hits against approximately 20,000 predefined track equations. To minimize latency, this operation is performed in parallel using combinatorial logic in FPGAs. Device resources are balanced by grouping track equations into four p_T bins, with each p_T bin corresponding to a single FPGA, as shown in Table 7.

Each track-finder FPGA outputs the six highest- p_T tracks it finds; these tracks are passed to a fifth FPGA for additional processing which includes sorting, matching tracks and CPS clusters, counting tracks, and calculating sector occupancy and total p_T . Tracks from each sector are sent over gigabit




Fig. 58. Transverse schematic view of a single 4.5° sector. A hypothetical track is overlaid on the eight CFT axial doublet layers and CPS axial layer. Track equations require a fiber hit on all eight CFT axial layers.

Table 7

Track equation distribution in track-finder FPGAs

p_T bin	p_T range		FPGA Resources
	(GeV/c)	Track equations	(system gates)
Maximum	> 10	3000	200k
High	$5 - 10$	3000	200k
Medium	$3 - 5$	4000	300k
Low	$1.5 - 3$	10000	500k

coaxial cables to L1Muon, where the tracks are matched to hits in the muon detector. Counts of tracks and occupancy data are passed downstream to the octant boards over LVDS links.

The next tier of DFE boards (CTOC) collects and sorts data within an octant (ten sectors). In L1 mode, these boards receive track counts and sector occupancy data from the DFEAs. The CTOC boards sum up the number of tracks, determine which sector had the most fibers hit, and check for isolated tracks. This information is passed on to a single DFE board (CTTT) where the individual trigger term bits are generated and sent to the trigger framework within $2.5 \mu\text{s}$ of the beam crossing. The CTTT can provide up to 96 trigger terms of which 55 are currently defined. Examples of CTTT trigger terms are: at least one or at least two tracks above a p_T threshold of 1.5, 3, 5, or 10 GeV/c or at least one isolated track above 5 or 10 GeV/c; and at least one track with a confirmation in the CPS for triggering on electrons.

The L1 TFW issues an L1Accept control bit to all readout electronics by considering trigger term bits from many different subsystems and makes a global L1 decision. Upon receiving an L1Accept, the AFE boards jump back a fixed number of events in their analog buffers and digitize the fiber data with 8-bit resolution. When the L1Accept control bit embedded in the




Fig. 59. CPS stereo subsystem hardware.




Fig. 60. FPS subsystem hardware.

data stream reaches the DFEA boards, they send a sorted list of tracks and CPS clusters to the CTOC boards. The lists are sorted by the CTOC and CTQD boards and passed downstream to preprocessor crates for more detailed analysis. These track lists, remapped onto the geometry of the silicon tracker, are also used as seeds for the L2STT. Additionally the CTOC and CTTT boards send copies of their input data to a L3 readout crate for monitoring and debugging.

- CPS stereo

The CPS stereo subsystem (Figure 59) provides information on the clusters in the two stereo layers of the CPS. Unlike the other CTT subsystems, CPS stereo does not generate L1 trigger terms. Rather, the DFE stereo (DFES) boards store discriminator bits and begin processing only after a L1 trigger decision has been made by the L1 TFW.

Event buffers in the DFES FPGAs store discriminator bits sent from the AFE boards. Upon receipt of the L1Accept control bit, the DFES boards extract the discriminator bits and search for clusters of hits in the stereo layers. Sorted lists of CPS clusters are sent on to the L2PS preprocessor for additional analysis. The DFES boards also send a copy of either their inputs (the discriminator bits) or their L2 outputs (the hit clusters) to the L3 readout system.

- FPS

The FPS subsystem (Figure 60) produces its own set of L1 trigger terms, which are passed to the L1 TFW. The overall structure of the FPS subsystem is similar to the CFT/CPS axial subsystem in that the FPS has three tiers of DFE boards: a finder (DFEF), a concentrator (FPSS), and a trigger term generator (FPTT).

The DFEF boards receive discriminator bits from the AFE boards and search for clusters of hits in the FPS fiber layers. A list of clusters is stored in the DFEF for later L2 readout while counts of clusters are passed downstream to the FPSS boards. The FPSS boards sum these cluster counts and

pass this information to the FPTT, where the L1 trigger terms are produced. These trigger terms must reach the TFW within $2.5 \mu\text{s}$ after beam crossing.

When the TFW issues an L1Accept, the FPS subsystem switches into readout mode. The AFE boards begin digitizing and reading out fiber data, while the DFEF boards extract cluster lists from their buffers. These lists of clusters are concatenated and sorted by the FPSS boards before being sent to the L2PS preprocessor. The FPSS and FPTT boards send their L1 input data to the L3 readout system.

- STT

The STT subsystem (Figure 57), described in detail in Section 9.3.5, is used to map the DFEA outputs onto the sixfold STT/SMT geometry. After an L1Accept has been issued, lists of CTT seed tracks generated by the DFEAs are concatenated and reformatted for the STT by a CTT subsystem consisting of two types of boards: six STOV (“overlay” or “overlap”) and six STSX (“sextant”) boards. The DFEAs have duplicate outputs: one set feeds the CTOCs, the other the STOVs (DFEA outputs shared between STT sextants) and STSXs (outputs unique to a sextant). The STOV outputs are fed into the STSXs; the STSXs send the data to L1 (lists of CTT tracks) via six optical fibers to the STT system.

9.1.4 Level 1 muon trigger

L1Muon looks for patterns consistent with muons using hits from muon wire chambers, muon scintillation counters, and tracks from the L1CTT. Field programmable gate arrays are used to perform combinatorial logic on roughly 60,000 muon channels and up to 480 tracks from L1CTT for every bunch crossing. Data from the detector front-ends are transmitted on custom Gbit/s serial links over standard coaxial cable. The serial link receivers and FPGAs are located on VME cards that reside in four custom VME crates on the detector platform.

The muon system (and L1Muon) is divided into central, north, and south regions. Each region is further divided into octants. Front-end data from each octant are processed by two L1Muon trigger cards (Figure 61). The scintillator trigger cards (MTC05) match central tracks to muon scintillator hits while the wire trigger cards (MTC10) match scintillator-confirmed track stubs in wire chambers between the three layers of the muon system. The octant decisions from each MTC05/MTC10 pair in a region are summed in the muon trigger crate managers (MTCMs) and sent to the muon trigger manager (MTM). The MTM forms 256 global L1Muon triggers and sends up to 32 of these to the L1 TFW. The download of the specific triggers is handled via EPICS software (Section 11). The total latency of the L1Muon trigger is about $3.20 \mu\text{s}$, driven by the central wire chambers (PDTs) and tracks from L1CTT.

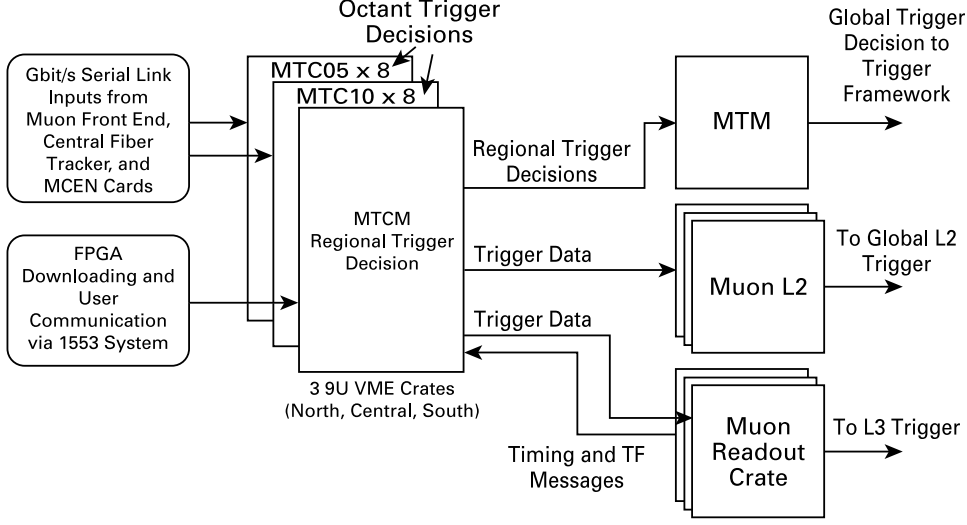


Fig. 61. Level 1 muon trigger system overview. Each octant has two trigger cards that process muon detector hit information and L1CTT tracks for that octant. The octant triggers for a given region are summed on the MTCM and sent to the MTM, which combines the three regions and sends up to 32 triggers to the TFW.

9.1.4.1 L1Muon hardware All detector inputs to L1Muon use Gbit/s serial links. The transmitters and receivers are 1.5" \times 2.2" daughter cards that are mounted on the muon front ends and on the L1Muon trigger cards. Each serial link can transmit up to $16 \times 7 = 112$ bits every bunch crossing. All MTC05, MTC10, and MTM trigger cards use a common motherboard with sixteen serial links and different flavor daughter cards that perform the MTC05, MTC10, and MTM logic.

The MTC05 cards match tracks from L1CTT to hits in the muon scintillator system. Each octant trigger card receives tracks from the L1CTT for the ten 4.5° sectors in that octant plus one sector of overlap on either side. Each sector sends the six highest- p_T tracks to L1Muon, and each track contains the CFT fiber position in the outer layer, p_T value, and sign of the track curvature in the central magnetic field. The triggers formed by the MTC05 cards include loose (track matched to A-layer scintillator) and tight (track matched to a scintillator road using the A and B layers) for four p_T thresholds (roughly 1.5, 3, 5, and 10 GeV/ c). Loose and tight scintillator-only triggers are also formed.

The MTC10 cards form triggers based on wire hits. In the central region, the hits from the wire chambers (PDTs) are sent directly to the trigger cards. The hits for each layer are used to form track stubs, or centroids, which are then used to confirm scintillator hits in each layer. Triggers are formed by matching centroid-confirmed scintillator hits between layers. In the forward region, the centroid finding is done by separate centroid-finding cards (MCENs), which subsequently send the centroids to the MTC10 cards. The MTC10 cards then use the centroid-confirmed scintillator hits to form loose (A-layer) and tight

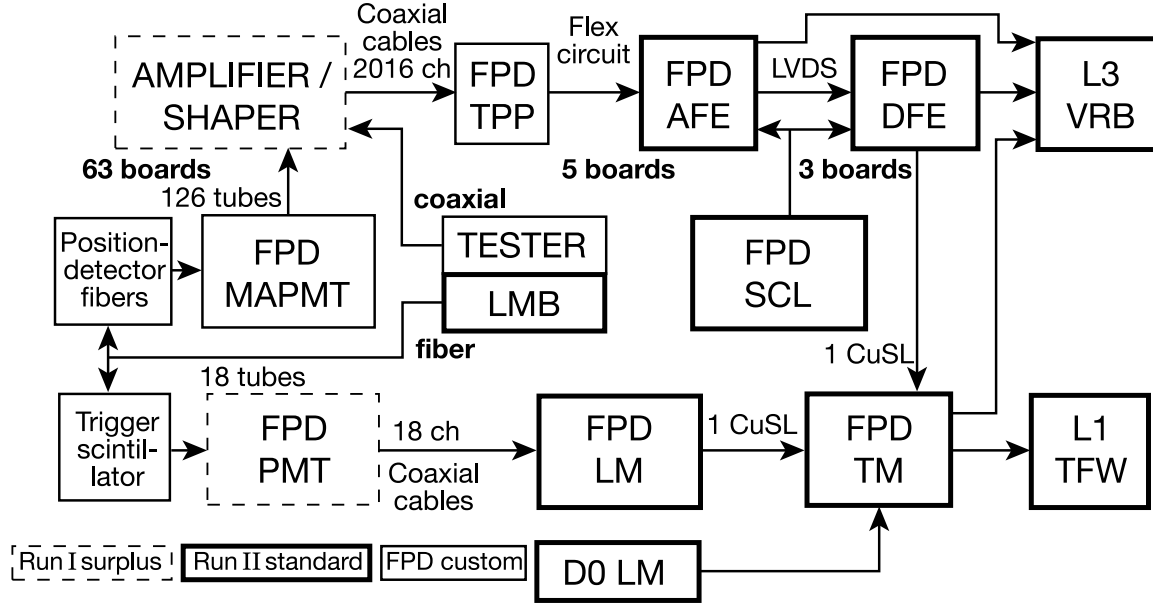


Fig. 62. A schematic view of the L1FPD trigger.

(A and B-layer) triggers.

The data from the various front end systems arrive asynchronously at L1Muon and must be synchronized before triggers can be formed for a given event. To accomplish this, all received data are written directly into FIFOs which are initially empty. When all FIFOs are not empty (i.e., they have all received data for the first bunch crossing), the data are read from the FIFOs and sent to the MTC05, MTC10, or MTM cards for trigger formation. In addition to synchronizing the data for a given event, the trigger cards also buffer the input data and trigger decisions pending global L1 and L2 trigger decisions. The input data and trigger decisions are stored in dual port memories and a pointer to the data is written into a FIFO. When an L1 or L2 Accept is received, the pointer is used to read the data for a particular event. The L1Muon trigger can also send all of the received input data from the detector front-ends to aid debugging.

9.1.5 Level 1 forward proton detector trigger

L1FPD selects events in which the outgoing beam particles pass through one or a combination of the nine FPD spectrometers. A schematic view of the L1FPD design is shown in Figure 62.

The signals produced in the FPD scintillating fibers are read out via 16-channel MAPMTs, amplified, and shaped before being passed to the AFE boards. The AFE boards provide analog fiber charge integration information to the L3 data acquisition system and pass discriminated signals via LVDS

connections to three DFE boards. At the DFE boards, pixel hit pattern recognition (a pixel is defined to be a u , x , and v fiber crossing) is used to compare the fiber hit information with valid detector hits stored in FPGAs. Three FPGAs containing between 600,000 and 1.5 million gates store generated L1 hit patterns for the nine spectrometers. The DFE boards also send their processed data to the L3 trigger system. The relative positions of the fibers in the different planes of each detector are used to define a finely segmented grid in which to search for hit patterns with optimal spatial resolution. Hit trigger conditions are set with different thresholds corresponding to different numbers of fiber layer hit coincidences.

The FPD trigger manager (TM) searches for coincidences between the discriminated hit signals of both position detectors of any FPD spectrometer. We use a set of nine FPD spectrometer (single diffractive) triggers with differing hit thresholds to select events in which at least one of the outgoing beam particles leaves the interaction region intact. Coincidences between the spectrometers on different sides of the interaction region are required to trigger on events in which both outgoing beam particles leave the interaction region intact in back-to-back configurations (elastic diffractive triggers) and in the larger set of other two-spectrometer configurations (double pomeron triggers). To reduce the contribution from beam halo spray particles that mimic a trigger signal, events with large hit multiplicities are rejected. The FPD information can be combined with timing information from the FPD trigger scintillator photomultiplier tubes (FPD LM), veto counters, and the luminosity monitor (DØ LM). The decision is then passed to the L1 TFW.

9.2 *The Level 2 trigger*

The L2 trigger provides detector-specific preprocessing engines and a global stage (L2Global) to test for correlations in physics signatures across detector subsystems. The L2 trigger system was designed to handle input rates of up to 10 kHz with a maximum accept rate of 1 kHz. L2 preprocessors collect data from the front-ends and L1 trigger processors and analyze these data to form physics objects. L2 can also combine data across detectors to form higher quality physics objects and examine event-wide correlations in all L2 physics objects. The L2Global processor selects events based on the set of 128 selections applied at L1 and additional script-controlled criteria. Events passing L2 are tagged for full readout and further analysis in the L3 trigger.

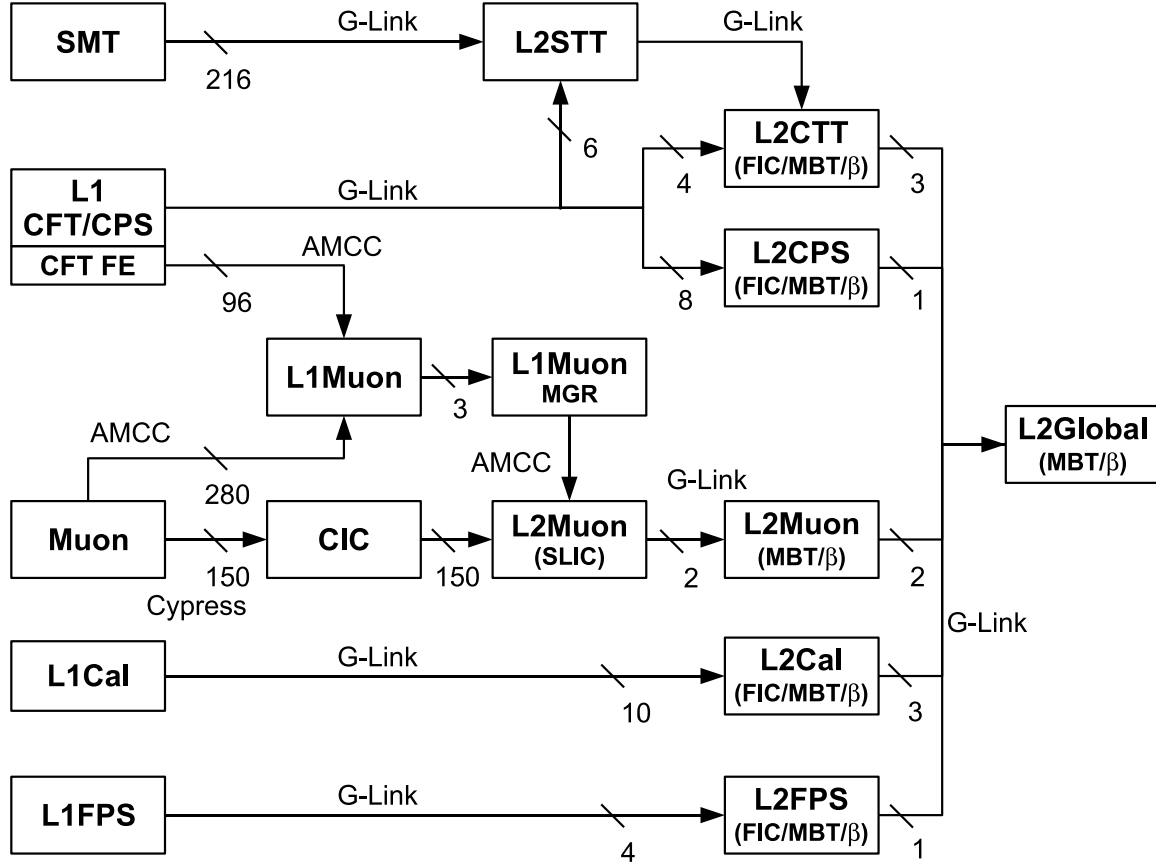


Fig. 63. L2 data paths and connections. Data paths are labeled according to transmission protocol, either 1.3 Gb/s fiber G-Links, 160 Mb/s copper Cypress Hotlinks, or 1.4 Gb/s copper AMCC links. The number of transmission cables is given for each pathway.

9.2.1 Architecture of the L2 trigger

L2 includes preprocessors for each detector subsystem and a global processor for integration of the data. Preprocessor subsystems include tracking, calorimeter, preshower, and muon systems. The subsystems work in parallel and trigger decisions are made in the L2Global stage based on physics objects reconstructed in the preprocessors. Preprocessing is performed either with serial CPU-based cards or with CPU cards plus highly parallelized DSP or programmable logic-based cards. The preprocessor and global stages function as 2-or-3-stage stochastic pipelines as illustrated in Figure 63.

Data arrive at the L2 system via two transmission protocols. Calorimeter and tracker data and signals from the L1 TFW are transmitted by 1.3 Gbit/s serial G-Links [103] on optical fibers. The muon system uses 160 Mbit/s Cypress Hotlink [104] transmitters on coaxial cables or standard CAT/6 cables, unshielded twisted pair (UTP) Hotlinks.




Fig. 64. Typical layout of cards in an L2 preprocessor crate. Cards that could be installed in parallel to increase input channels or processing power are designated by stacked boxes.

9.2.2 L2 crates

The L2 system occupies 9U VME crates built to the VME64/VIPA [105] standard. In addition to a 64-bit VME bus, these crates are instrumented with a custom high-bandwidth 128-bit bus (Magic Bus [106] or Mbus) for fast intra-crate data flow and communication. The Mbus supports data rates of up to 320 Mbit/s. The VME backplane is used to read out events accepted by L2 into L3 and for control and monitoring operations.

A typical L2 preprocessor/global crate contains the following devices:

- VME controller and dual-port-memory (DPM) card used for downloading runtime parameters and for reporting monitoring data to the TCC
- a VME single board computer (SBC) used to send data to L3
- one or more Mbus transceiver (MBT) cards (Hotlink-to-Mbus interface, Section 9.2.3)
- one or more sets of fiber converter (FIC) and VME transition (VTM) cards (fiber-to-hotlink interface, Section 9.2.3)
- one or more L2 β processor cards (VME and Mbus hosted single board computers, Section 9.2.4) for data processing.

The layout of the preprocessor/global crates is shown in Figure 64. L2STT (Section 9.3.5) and L2Muon (Section 9.3.2) use additional specialized cards for processing their data before sending track information to the L2CTT and L2Muon preprocessors respectively.

9.2.3 Magic bus transceiver and data conversion cards

Data arriving at L2 must first be converted into the correct physical format before processing can begin. The coaxial Hotlink data from the muon system is received by analog cable input converter (CIC) cards. These cards retransmit the data to the first stage of L2Muon as differential UTP Hotlink signals. A relative of the CIC card, the serial fanout (SFO), selectively fans out 1 to 12 UTP Hotlinks where signal duplication is necessary. Similarly, FIC cards receive G-Link data and buffer up to sixteen events while converting the inputs to differential UTP Hotlink signals. The incoming fibers are received by a VME

transition module (VTM) located at the rear of the crate.

Data are passed to L2 β processors via the Mbus. An MBT card transfers data to and from the Mbus. The MBT accepts up to seven Hotlink inputs, assembles the data into events, and broadcasts them on the Mbus backplane to the L2 β cards. The high speed, 128-bit input path is equipped with a buffer for up to sixteen events. There are two Hotlink output paths that are targets of the 128-bit Mbus programmed I/O on the L2 β cards. These are used by preprocessors to send their output to L2Global.

One MBT, designated the “pilot,” coordinates the event broadcast across all MBTs and receives SCL information for the crate from the TFW, including the L1Accept information, L2 decision information, and SCL initialize messages. The pilot MBT in the L2 global crate transmits the L2 decision to the TFW over sixteen PECL output channels. A demultiplexing card receives eight 16-bit words and transmits the 128-bit decision to the TFW.

9.2.4 L2 β processors

L2 β processor cards are used for analysis and control of data flow in each preprocessor subsystem crate. The L2 β processors replaced the original L2Alpha boards¹ early in Run II and are composed of commercially produced (Compact PCI standard [109]) SBCs mounted on 6U-to-9U VME adapter cards. Each SBC provides dual 1 GHz Pentium processors. The SBC connects to the adapter via a 64-bit, 33/66 MHz PCI bridge. The 9U adapter, controlled by the SBC, implements all DØ-specific protocols for Mbus and TFW connections. Custom I/O functions on this card are implemented in a single FPGA (Xilinx XCV405E) plus assorted logic converters and drivers. The FPGA is used to implement internal data FIFOs and address translation tables for broadcasting data from the Mbus to CPU memory, reducing the complexity of the adapter card. Mbus programmed I/O and various other trigger system interfaces, via front panel connections and VME user-defined pins, are also implemented. The adaptor also provides a 64-bit PCI-to-VME interface via a Tundra Universe II chip [110].

The L2 β processors run a GNU/Linux system and all programs are written in C++. Real time trigger performance is achieved by restricting system calls (including dynamic memory allocation) and promoting the L2 executable to a non-

¹ L2Alpha processors were similar to those constructed by the CDF experiment [107,108]. They were based on the PC164 motherboard manufactured by Digital Semiconductor (now Compaq). Each card contained a 500 MHz 21164 64-bit Digital Alpha CPU and 21164 PCI interface integrated with the necessary hardware interfaces. The L2Alphas have been supplanted by the higher performance and upgradable L2 β cards.

interruptible scheduling queue. The second CPU provides full access to the SBC for runtime monitoring and control.

9.3 *L2 preprocessor subsystems*

9.3.1 *L2Cal*

The calorimeter preprocessor system identifies jets and electrons/photons and calculates the \cancel{E}_T for the global processor. The worker code may be combined to run serially in a single processor or placed in separate processor cards to increase throughput as data rates and detector occupancies grow with luminosity. Each processor uses the E_T data from the 2560 calorimeter trigger towers. The input data arrives on ten input links which together transport 3 kB/event of tower transverse energy data, including both EM towers and the EM+H tower sums.

9.3.1.1 Jet algorithm The jet algorithm operates by clustering $n \times n$ (currently $n = 5$) groups of calorimeter trigger towers which are centered on seed towers. The seed towers are E_T -ordered with $E_T \geq 2$ GeV. Overlapping candidates may be reported as separate jets depending on programmable criteria based on the number of shared towers (otherwise the highest- E_T jet of the overlapping candidates is kept). The list of jets is passed to L2Global which applies jet cuts as defined by the trigger menu.

9.3.1.2 Electron/photon algorithm The electron/photon algorithm begins by defining an E_T -ordered list of EM towers with E_T above 1 GeV. For each seed tower, the neighboring tower with the largest E_T is combined with the seed to make an EM cluster. The EM energy fraction of the leading- and sub-leading- E_T trigger towers of the cluster and the amount of total E_T in a 3×3 tower array surrounding the seed tower of the cluster are used to reduce background. The final list of electron candidates is sent to L2Global to apply the trigger requirements.

9.3.1.3 Missing transverse energy algorithm The L2 calorimeter \cancel{E}_T algorithm calculates the vector sum \cancel{E}_T from the individual trigger tower total- E_T energies passed from L1. It is capable of calculating the \cancel{E}_T for different minimum tower E_T s and η ranges.




Fig. 65. Layout of cards in an L2Muon crate. Two stages of processing are completed in a single crate using SLICs and L2 β processors. Central and forward muon regions are processed in separate crates of similar configuration.

9.3.2 *L2Muon*

L2Muon uses calibration and more precise timing information to improve the quality of the muon candidates [111]. It receives the L1Muon output and approximately 150 front-end data sources (from the PDTs and the scintillation counters). The muon candidates contain the track p_T , the η and ϕ coordinates of the muon, and quality and timing information of the muon component stubs.

L2Muon implements one extra level of preprocessing in the stochastic pipeline sequence. The first L2Muon stage incorporates 80 Texas Instruments 200-MHz DSPs in a parallel processing scheme. Each DSP is responsible for finding track stubs in a small region of the detector, and therefore, the execution time of the algorithms is independent of the number of hits. The DSPs are geographically organized in eleven central and five forward 9U VME boards (second level input computers or SLICs), with each SLIC containing four worker and one administrator DSP chips. The SLICs are programmable in C. The layout of an L2Muon crate is shown in Figure 65. The DSP algorithms, characterized by the detector region (central or forward) and the input or sub-detector plane (L1, A or BC muon layers), make use of detector symmetry to run the same basic processing code. The MBT sends the stubs found by the SLICs to the L2 β processor. The L2 β board uses the track segments to construct integrated muon candidates with an associated p_T and quality and timing masks.

9.3.3 *L2PS*

Both the central and the forward preshower detectors are designed to provide high electron detection efficiency, electron-photon separation and high background (charged hadron) rejection at the trigger level. This is accomplished by providing evidence for early shower development and by giving a good spatial point for comparison with calorimeter clusters or tracks. At L2, the CPS and FPS are treated as separate detectors and their data are processed

independently.

Upon L1Accept, CPS axial clusters from each of the eighty azimuthal trigger sectors are combined into quadrants in azimuth and transmitted to the L2 preprocessor. Stereo clusters are sent directly to the L2 preshower preprocessor. Figure 63 shows the L1 to L2 data transfer path into the preshower subsystem. The L2CPS preprocessor receives axial (x) clusters over four G-Links, each serving one azimuthal quadrant. The CPS stereo data are transmitted over four G-Links, two for each hemisphere (positive/negative η), including both u and v layer clusters. Axial clusters are tagged by L1 with the presence or absence of a CFT track. Each of the G-Links can carry a maximum of 48 axial CPS clusters or 96 stereo clusters.

In the L2 β processor, the CPS cluster centroids are compared to produce η, ϕ coordinates for clusters that match in three layers. The presence or absence of CFT trigger tracks associated with CPS x clusters is also provided, and output clusters are flagged as electrons (when there is a track associated with the cluster) or photons (no track). The η, ϕ coordinates are binned to correspond to the calorimeter trigger tower geometry of $\eta \times \phi = 0.2 \times 0.2$. A window of width 0.05 is drawn around each calorimeter trigger tower, and any preshower hit in this η, ϕ region designated a calorimeter match. The FPS provides similar functionality and is the only source of forward tracking information available before the L3 trigger.

9.3.4 L2CTT

The L2CTT preprocessor takes inputs from the L1CFT and the L2STT. This preprocessor system has been designed to operate in two different modes: *i*) with input tracks straight from L1CFT and *ii*) with input tracks from L2STT which receives input from the L1CFT and SMT barrels.

In the first mode of operation, L2CTT reads in the track lists from different ϕ regions of the L1 tracking trigger system and concatenates these into a single p_T -sorted list. The p_T measurements are refined using additional hit and tracking information provided by the L1 readout. For each track, the azimuthal angle with respect to the beamline, ϕ_0 , is determined. The value of the azimuthal angle at the third layer of the EM calorimeter, ϕ_{em3} , is also calculated (ϕ_{em3} is different from ϕ_0 due to the bending of tracks in the solenoidal magnetic field). Finally each track is evaluated according to several isolation criteria to enhance the trigger capabilities for τ leptons. The p_T -sorted list of L2 tracks is reported to L2Global. In the second mode of operation, input data are provided by the STT along with refined L2 track p_T s. Only ϕ_0 , ϕ_{em3} , and isolation are calculated for these data. However two separate lists of L2 tracks are passed on to L2Global, one sorted according to p_T and another




Fig. 66. The definition of roads based on L1 tracks and SMT hit selection in L2STT. sorted according to impact parameter.

9.3.5 Level 2 silicon track trigger

The L2STT performs online pattern recognition in the data from the SMT. It reconstructs charged particle tracks found in the CFT at L1 with increased precision by utilizing the much finer spatial resolution of the SMT. The L2STT provides better momentum resolution for the trigger and the ability to trigger on events containing tracks with large impact parameters.

The L2STT improves the momentum measurement of charged particle tracks at the trigger level. Requiring hits in the SMT helps reject spurious L1 triggers from accidental track patterns in the CFT. The primary physics justification of the L2STT is its ability to measure the impact parameter of tracks precisely enough to tag the decays of long-lived particles, specifically B hadrons.

9.3.5.1 Conceptual design Figure 66 shows the basic principle of the L2STT. For each event, the L1CTT sends a list of tracks to the L2STT. A road is defined around each track and the SMT hits within the road are associated with the track. The L2STT uses only the hits in the axial strips of the silicon ladders, which define points in the $r - \phi$ plane. The L2STT uses the hits in the innermost and outermost layers of the CFT and hits in at least three of the four layers of the SMT to fit the track parameters. The results of the fits are sent to L2Global.

The SMT barrel ladders are arranged in twelve sectors, each covering 30° in azimuth. The ladders of adjacent sectors overlap slightly such that more than 98% of all tracks are contained in a single sector. The L2STT therefore treats all 30° sectors independently.

9.3.5.2 System architecture The L2STT consists mainly of custom designed digital electronics modules. All custom modules are designed to plug

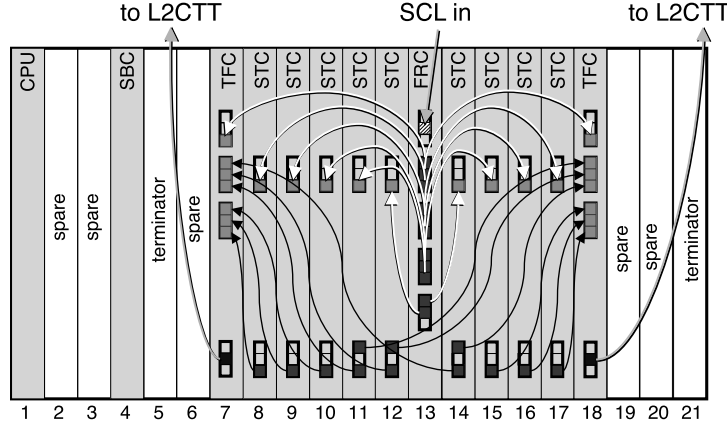


Fig. 67. Layout of an L2STT crate.

into a motherboard, and a common motherboard design is used throughout the system. Data input from the SMT detector and the L1CTT is via optical fiber serial links into receiver cards (VTMs) located in the rear card cage of the VME64/VIPA crates that house the L2STT electronics. The data are processed in large FPGAs and/or DSPs on a logic daughterboard that sits on the motherboard. There are three different types of such daughterboards in the system. Data are communicated between modules in an L2STT crate using serial link transmitter and receiver cards. Each module is equipped with a daughterboard that buffers data for readout through the data acquisition system once an event has been accepted by the trigger system. The logic daughterboard is connected to the buffer cards, the serial links, and to the VME backplane by three PCI busses on the motherboard. The VME bus is used for initialization and monitoring and to read data out of the buffer cards.

The three types of logic daughterboards are the fiber road card (FRC), silicon trigger card (STC), and track fit card (TFC). The FRC receives data from the L1CTT and the trigger framework which it fans out to the other modules. The FRC also manages the storage of data in the buffer cards. The STC receives the SMT data, clusters hits in adjacent strips, and associates SMT clusters with roads. The TFC performs the final hit selection and fits the tracks.

L2STT consists of six identical VIPA crates, each serving two 30° degree sectors. Each crate is equipped with one FRC module, nine STC modules, and two TFC modules. Each crate also has a CPU board to program the FPGAs and DSPs, download constants, and monitor system performance. An SBC is used to read data out of the buffer cards and feed them into the data acquisition system. Figure 67 shows the layout of an L2STT crate.

Each of the twelve TFCs provides a list of tracks from its 30° sector. The track information includes track parameters and the χ^2 of the fit as well as additional information about the cluster selection and fit. One TFC in each




Fig. 68. Impact parameter resolution as a function of p_T determined using simulated single muon events.

of the six crates also transmits a list of the initial L1CTT tracks. These data are transmitted to the L2CTT, where the tracks are sorted by p_T and impact parameter and passed to L2Global to be used in the trigger decision. The data are also sent to a buffer card for readout to L3.

9.3.5.3 L2STT performance The L2STT helps to select events with an enhanced b -flavor content by measuring the impact parameter b of reconstructed tracks with respect to the proton and antiproton beams. Figure 68 shows the impact parameter resolution σ_b obtained from simulated single muon events. The impact parameter resolution has a p_T dependence introduced by multiple scattering. In the trigger, the effect of this p_T dependence can be reduced by using the impact parameter significance $S_b \equiv b/\sigma_b$ instead of the impact parameter b . The uncertainty $\sigma_b = \sqrt{(18 \mu\text{m})^2 + [(50 \mu\text{m GeV}/c)/p_T]^2}$ (for tracks with four SMT clusters in the fit) takes into account the effect of multiple scattering.

9.4 L2Global

L2Global is the first level of the trigger to examine correlations across all detector systems. The L2Global worker is responsible for making trigger decisions based on the objects identified by the L2 preprocessors. Trigger decisions are made by creating global physics objects. These objects can be based directly on the objects reported by the preprocessors or can be created by combining objects from different preprocessors. The L2Global worker imposes cuts on the global physics objects according to configuration information it receives from the TCC based on the downloaded trigger menu.

9.4.1 Inputs to L2Global

After an L1Accept is issued, L1 sends a trigger decision mask to L2Global. The SCL sends notification of an L1 accepted event to every geographic section that needs to read out to L3. In the case of L2 preprocessors, the receipt of the L1Accept SCL message means that the preprocessor must send at least a header block to L2Global for this event, and must prepare at least a header block for eventual L3 readout in case the event passes L2. When L2Global receives the L1Accept, L2Global must perform a decision cycle and prepare at least a header block for L3 readout if the event passes (or if it is an Unbiased Sample as described below).

For each event, the L2Global worker uses the L1 information to decide which algorithms to run on the data from the preprocessors. It then makes a trigger decision, and returns this decision to the trigger framework. The trigger list specifies which trigger conditions L2 will impose for each run. The trigger list can change as frequently as every run, but attempts are made to achieve stability in the trigger list to reduce complexity at the analysis level. The trigger list is downloaded to the L2Global crate by the TCC, which receives its instructions from COOR.

A number of L1 qualifier bits may be sent along with a given L1Accept. Of these L1 qualifiers, the ones of greatest interest to L2Global are *Unbiased Sample* (UBS), *Forced Write*, and *Collect Status*. For a UBS event, L2Global sends the event to L3 whether or not any of the L2 bits actually passed, marking the condition in the L2Global event header. The actual L2 decision is also recorded. Such events occur at a rate determined by the trigger programming, for an independently adjustable fraction of events passing each L1 bit. A secondary effect of the UBS qualifier in L2Global is that additional information is written to L3 to assist in debugging the event. In particular, the output from L2Global is expanded to include more information to allow detailed checking of the processing. Events marked with the UBS qualifier are routed to a special data stream for system monitoring.

The *Forced Write* qualifier provides a mechanism to test new trigger definitions. Unlike the UBS qualifier, *Forced Write* produces detailed output at every firing of a L1 trigger bit that is marked with this qualifier. The *Forced Write* qualifier has the same effect on L2 as the UBS qualifier, however *Forced Write* events are routed to the standard output stream for offline analysis.

After receipt of an event with a *Collect Status* qualifier, L2Global (and all preprocessors) capture internal scaler and other monitoring information (including distributions of processing times and pipeline occupancies) for readout by the TCC. The monitoring blocks are tagged with the L1 crossing number of the event with the *Collect Status* qualifier, so the TCC can assemble a con-

sistent set of statistics in the L1 and L2 systems. *Collect Status* qualifiers are generated approximately once every five seconds in a typical data run.

9.4.2 *Trigger scripts*

The L2Global worker begins every event by checking which of the L1 trigger bits fired. The trigger list specifies which L2 script is associated with each L1 bit. The L2 script is the trigger condition that must be satisfied in order for the L2 trigger to fire for a given L1 bit. The L2 script is specified by a number of filters and a minimum number of objects required to pass each filter. An example of a script is an electromagnetic object filter and a minimum of two objects. This script passes if there are two electromagnetic objects in the event that satisfied the conditions of the filter. If any script passes, the event passes L2 and is processed by the L3 trigger.

9.4.3 *Tools and filters*

Tools and filters provide the main functions of the L2Global worker. The filters make up the scripts described above and in turn rely on tools. Tools are C++ classes which build a specific type of L2Global object. L2Global objects are based on preprocessor objects. A tool starts with a list of preprocessor objects and applies selection criteria to decide which preprocessor objects should be made into global objects. A tool can also correlate information from two preprocessors by combining two preprocessor objects to form one global object. An example of this is the combination of a track object and an EM object that come from different preprocessors, but refer to the same electron candidate. The tools are written to be flexible and are configurable through the trigger list. For the above example, one trigger list parameter specifies whether or not EM objects from the preprocessor should be matched to tracks from the tracking system.

The tools produce lists of global objects. The filters then use these lists of global objects to make a trigger decision by imposing trigger requirements on the objects. For the case of an EM object, the electromagnetic fraction, transverse momentum, and isolation can be required to have values above or below specified thresholds to pass the filter. The filter generates its own list of objects that satisfied the trigger criteria. Tools and filters can also operate on global objects passing other filters to construct global objects of greater complexity. At the conclusion of processing, the script checks to see if there are at least the minimum number of objects required to pass the script.

9.5 *Level 3 trigger*

The L3 trigger provides essential additional rejection both to enrich the physics samples and to maintain an acceptable throughput which can be recorded to tape. A high level, fully programmable software trigger, L3 performs a limited reconstruction of the event, reducing a nominal 1 kHz input rate to 50 Hz for data recorded for offline analysis. Its decisions are based on complete physics objects as well as on the relationships between such objects (such as the rapidity or azimuthal angle separating physics objects or their invariant mass). Candidate physics objects, or relations between them, are generated by object-specific software algorithms (filter tools). Tools perform the bulk of the work: unpacking raw data, locating hits, forming clusters, applying calibration, and reconstructing electrons, muons, taus, jets, vertices, and \cancel{E}_T . Reference sets (refsets) of programmable algorithm parameters are input to the tools via the trigger programming triggerlist. The refsets define the physics objects precisely (jet refsets specify cone size, for example, and electron refsets, the electromagnetic fraction, among other characteristics) for each invocation of the filter tool. All tools cache their results to expedite possible multiple calls within the same event, and if the event is accepted, add to the data block position and kinematic quantities.

Individual calls to the tools are made by filters that define the specific selection criteria employed by a tool or imposed on its results. These criteria include the refset used by the tool, as well as thresholds and other simple cuts applied by the filter on a tool's results (for example, the requirement of two jets within a given pseudorapidity range above a fixed E_T threshold). Filter results are keyed for access by other filters, so in addition a key to a previous filter's results can be included in these parameter sets. Part of the trigger programming, this information can be changed with each download.

The triggerlist programming includes blocks of filter scripts, that specify one or more filters, and which define the L3 trigger condition for each L3 trigger or filter bit. Each L3 filter script is associated with a L2 bit; multiple L3 scripts may be associated with each L2 bit. Failure to pass an individual filter terminates execution of the script, calling no further tools, and skipping to the next filter bit's script. Only when all filters in a script are satisfied, is the trigger satisfied and the event sent to the host cluster to be recorded.

9.5.1 *ScriptRunner*

Each filter tool receives event data under the direction of the L3 framework's interface to the tools, ScriptRunner. At the beginning of a run, ScriptRunner initiates parsing of the tool refsets and the filter scripts. L3-relevant pieces of

the triggerlist (and event data) are its input. Needed tools individually initialize themselves with the necessary calibration constants and resets. ScriptRunner then processes any errors and parses the scripts to build the execution tree. ScriptRunner's output is filtered event data dispatched to the data logger, monitoring information (individual and combined filter rates and timing of both filter scripts and individual tools collected in the event loop), and error messages. At the end of a run (and upon request) ScriptRunner extracts and reports monitor information and rates.

In the order specified by the trigger list, the execution path proceeds along the branches of each L2 trigger bit that has been set. Each filter branch is traversed, with tools called in the given filter script order until one fails or all pass; execution then returns to the filter bit level and proceeds to a sister filter bit branch, if one exists, or continues to the next trigger bit. Control is returned to ScriptRunner as soon as a filter fails or after all filters have passed. Figure 69 shows an example of the sequence within one possible filter script branch of the execution tree.

9.5.2 Available physics object tools

9.5.2.1 Level 3 jets and electrons The L3 jet tool relies on the high-precision calorimeter readout and primary vertex position available at L3 and the improved energy and position resolution this information makes possible. Implementing a simple cone algorithm and performing a removal of hot calorimeter cells, the L3 jet tool is able to sharpen the turn-on curve dramatically. Applying offline cleanup cuts on collected data shows sharp turn-ons to nearly 100% efficiency for both jet and electron filters (see Figure 70). Rejection factors of 20–50 have been realized for the various jet triggers.

A simple 0.25 jet cone, with cuts on E_T , electromagnetic fraction (> 0.9), and transverse shower shape, forms the basic electron tool. An electron candidate match to a preshower detector signal can also be required. Combining contiguous preshower strips above threshold into single-layer-clusters (SLCs) for each layer and hemisphere, the L3 preshower tool applies an SLC energy threshold cut and then searches for geometrically allowed combinations (of one SLC from each layer), matched within position errors and with a reasonable energy correlation. Studies with MC $Z \rightarrow e^+e^-$ samples overlaid with six minimum bias events show the CPS tool provides 99% efficiency in under 10 msec/event.

9.5.2.2 Level 3 muons Wire and scintillator hits are used in the reconstruction of muon track segments inside and outside the toroid. Track-finding links these segments to identify muon tracks in three dimensions. Resets de-

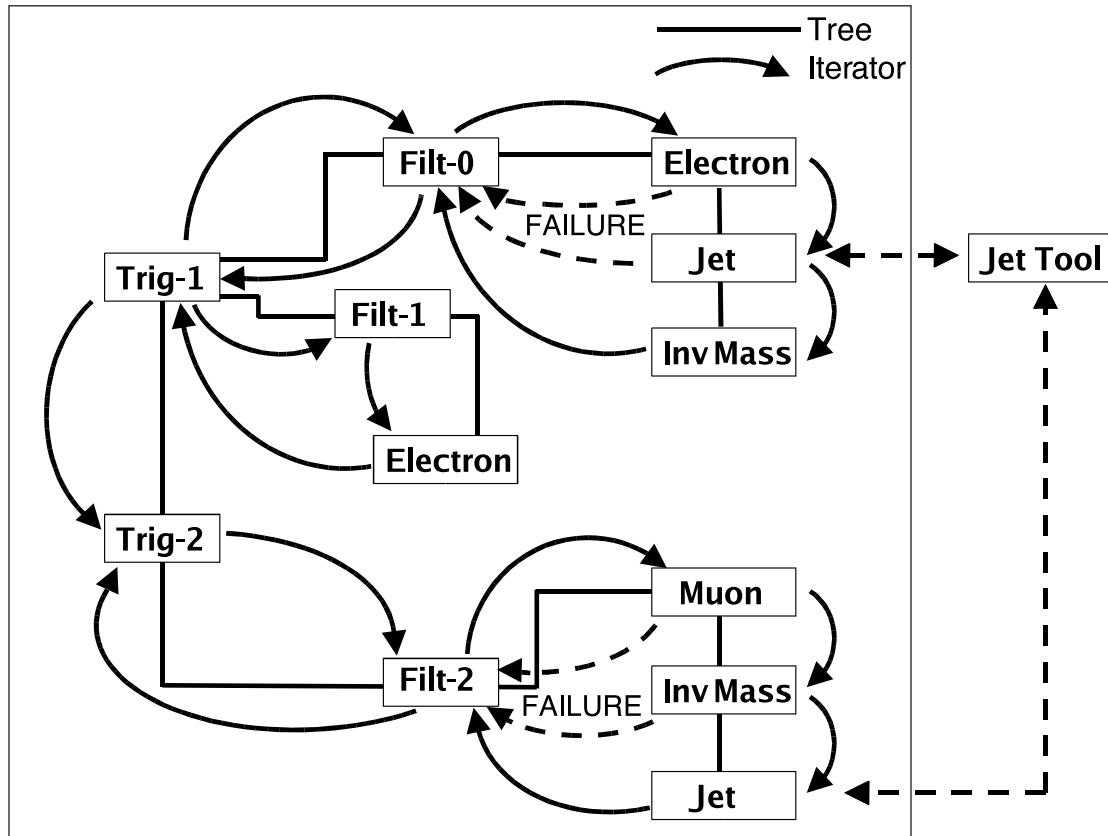


Fig. 69. The execution tree. L3 filter bit scripts branch from each L2 trigger bit, with a chain of individual filters hanging from each. Execution returns from the filters (on failure or completion of the entire branch) to the filter bit script level. After all filter bit scripts have been traversed, control continues to the next valid trigger bit branch.




Fig. 70. The Level 3 jet trigger turn-on. Efficiency is plotted as a function of offline leading jet p_T . A selection cut of 15 GeV/c has been applied by the L3 filter. Statistical errors are included, though small enough to be obscured by the symbol in most instances.

termine the minimum number of hits in each detector layer required for valid tracks. In addition to unpacking data and determining tracks in the geographic regions defined by L2, the L3 muon tool can call subordinate tools that utilize inner tracking and calorimetry.

L3 improves on L2's ability to separate prompt, slow, and out-of-time hits by fitting the available scintillator hits along a track to the particle's velocity. Remnant cosmic ray muons are recognized by the out-of-time signals in addition to their association with a penetrating track opposite their candidate track.

The availability of vertex constraints and the ability to match central tracks makes possible improved momentum resolution in L3 compared with L2. A track match tool extrapolates muon tracks to the central tracker, defining a region to be searched. If more than one track is available, the track that minimizes the χ^2 calculated using the angular distances in θ between the central and extrapolated muon tracks is selected. Additionally, matching local muon tracks to paths of minimum ionizing particles in the calorimeter distinguishes isolated from non-isolated muons.

9.5.2.3 Missing transverse energy The L3 \cancel{E}_T tool is based on calorimeter cells. Using corrected calorimeter energy (assuming a nominal (0,0,0) vertex), the \cancel{E}_T is calculated through intermediate pseudorapidity ring sums (allowing quick recalculation by geometrically corrected ring sums when the primary vertex tool returns a vertex position). The tool calculates \cancel{E}_T , ϕ_{MET} , total scalar E_T , and a \cancel{E}_T significance.

The filters that make up the L3 trigger can be flexible and complete. For example, adding a \cancel{E}_T filter to loose electron or muon filters, in parallel to single lepton scripts, provides redundancy for triggering on W boson events; adding \cancel{E}_T to lepton + jets top quark filters provides additional rejection factors (a factor of three for electron and eight for muon triggers).

9.5.2.4 Level 3 Tracking CFT track-finding can be performed either in specified regions or across the entire detector. At initialization, the calculated x, y, ϕ fiber positions are stored in lookup tables. Adjacent hits are merged into clusters with their x, y, ϕ position averaged. Tracking is first done using a fast circle fit through the axial layers, identifying candidates as arcs through the origin and defining an arc length S in the $x - y$ plane. A straight line fit in the $S - z$ plane then determines the track helix parameters.

A link-and-tree algorithm joins clustered hits from different layers. Starting from a link in the outer layer (and continuing recursively), candidate tracks are

built by adding links from adjacent layers, extending the path length as long as the curvature is consistent with the preceding link. The longest extended path found starting from the initial link is kept as a track candidate. The radius defined by the circle through clusters (and the origin) must exceed that corresponding to minimum trigger p_T .

SMT tracking modifies this method: segments connect neighboring hits between points within a specified $\Delta\phi$. Segment paths are linked when their z -slope and $\Delta\phi$ are within (tunable) specifications. Unless seeded by a location set by an earlier tool's candidates, the algorithm begins with the outermost SMT layer, looking for the longest paths toward the center, fitting these to a helix. The path with the smallest χ^2 is selected.

A CFT-track-based vertexing tool offers substantially higher efficiency and purity than an SMT-hit-based algorithm. Its use sharpens the turn-on curve for jet triggers, and provides rejection against events whose primary vertex is at large $|z|$. Monte Carlo studies suggest that the tool achieves 500 μm vertex resolution with a four-fold rejection with no loss in efficiency. The CFT-track-based vertex is used for jet identification and the \cancel{E}_T calculation.

A global (SMT plus CFT) high-momentum-track finder starts from axial CFT tracks (with matched stereo clusters) propagated into the SMT. If the CFT axial/stereo match fails, the CFT-SMT match is done in x, y only. A stand-alone global track filter is a useful addition to single muon and electron triggers. Starting from axial CFT tracks (with matched stereo clusters) propagated into the SMT, the L3 tracker makes an independent selection of individual charged high- p_T tracks. This algorithm (with 60 μm distance-of-closest-approach (DCA) resolution for central tracks) runs in less than 200 ms. Studies of off-line track-matched electrons in a $Z \rightarrow e^+e^-$ data sample show a 60% overall efficiency (95% within the CFT-axial acceptance). With a one-track requirement, the central di-electron trigger has a rejection of 150. A L3 single electron filter (plus shower shape and one-track cuts) enjoys a rejection of 300. For single muons, requiring one track gives a rejection of 45. By requiring at least two SMT hits plus stereo information, the number of fake tracks can be reduced by a factor of two.

The use of tracking information in the L3 filter reflects a strategy of parallel single-electron filters which increase efficiency at low E_T and still provide redundancy at high E_T (for cross checks). These parallel triggers allow high- E_T filters with loose cuts (where rejection is not critical) by introducing tighter cuts at lower E_T . For example, a combined rejection of 27 was achieved by moving the L3 electron filter threshold to 30 GeV and introducing a 20-GeV-threshold filter with shower shape cuts. Online studies confirm the individual rejections predicted offline, and show the total rejection (OR of five single electron filters) is 14. Similarly L3 global track filters (run on muon triggers)

are complemented by a suite of L3 muon filters running the stand-alone local muon filter for high- p_T single muons. A J/ψ sample with both muons matched to central tracks offline shows L3 gives an efficiency of 80% per muon.

Online monitoring keeps track of current beam spot information (the mean position and spread in x, y along with tilts in x, z and y, z). L3 obtains this information at the beginning of each run, allowing L3 to calculate a fully 3-dimensional primary vertex for each event. By recalculating L3 track parameters using the 3-d vertex, L3 is capable of triggering on the impact parameter of tracks. The DCA resolution with respect to the primary vertex is $25\ \mu\text{m}$. With input provided by a tracking, jet and vertex tool, event, jet and track b -quark-probabilities based on the signed impact parameters of tracks can be calculated, and b -tagging implemented in L3. Such a requirement added to a track-matched muon filter has been demonstrated to halve the pass rate for this trigger with no loss of (offline) double-tagged b jets.

9.5.2.5 Relational filters Additional higher level selections, made on the relationship between physics tool candidates, are implemented at the filter level. Examples include the invariant mass filter, an acoplanarity tool that selects events with the two leading- p_T jet candidates separated by a polar angle between ϕ_{min} and ϕ_{max} , and an H_T tool that applies a selection cut to the scalar sum of all jet filter candidate E_T 's.

9.5.3 *Special filter tools*

For monitoring purposes, events can be written out regardless of the L3 trigger decision. Special “mark and pass” runs record every L2-accepted event and mark them with the results of all filter tools of every script under the L1/L2 bit(s) satisfied. Additionally, a “pass one of n ” option for each script independently specifies a fraction of events forced through marked but unfiltered. Every event selected in this way passes into a special “monitor” stream, distinct from the physics stream. Monitor stream data are collected continuously during regular data-taking.

When necessary to prescale events, L3 can mark a filter bit as failed without running the filter script. The L3Prescale filter's implementation uses selection by random number generation (seeded uniquely by node) so that collectively, successive short runs still see events in the correct prescale fraction. This implementation by random numbers initialized differently in each node is the same means by which “pass 1 of n ” events are generated.

9.5.4 Online monitoring

Online monitoring of the L3 system is done in the control room during regular data-taking. Quick analysis of the event record from randomly sampled online events produces distributions of physics quantities (for failed as well as passed candidate events). Plots of electron, jet, muon, tau, and global track multiplicity, E_T , ϕ and η are continuously reviewed. In addition, a comparator package looks for discrepancies between online quantities and those computed when the data are run through the trigger simulator offline.

10 Data Acquisition System

The DØ data acquisition (DAQ) system consists of the L3DAQ and online host systems. L3DAQ transports detector component data located in VME readout crates to the processing nodes of the L3 trigger filtering farm; the online host receives data from the L3 farm nodes and distribut

10.1 L3DAQ

The L3DAQ system's designed bandwidth is 250 Mbyte/s, corresponding to an average event size of 250 kbyte at a L2 trigger accept rate of 1 kHz. As shown in Figure 71, the system is built around a single Cisco 6509 [112] ethernet switch. A schematic of the communication and data flow in the system is shown in Figure 72. All nodes in the system are based on commodity computers and run the Linux operating system. TCP/IP sockets implemented via the ACE [113] C++ network and utility library are used for all communication and data transfers.

Sixty-three VME crates must be read out for each event, each containing 1–20 kB of data distributed among VME modules. An SBC in each VME crate reads out the VME modules and sends the data to one or more farm nodes specified by routing instructions received from the routing master (RM) process. An event builder (EVB) process on each farm node builds a complete event from the event fragments and makes it available to L3 trigger filter processes.

The supervisor process provides the interface between the main DØ run control (COOR) and the L3DAQ system. When a new run is configured, the supervisor passes run and general trigger information to the RM and passes the COOR-provided L3 filter configuration to the IO/EVB process on relevant farm nodes, where it is cached and passed on to the L3 filter processes.

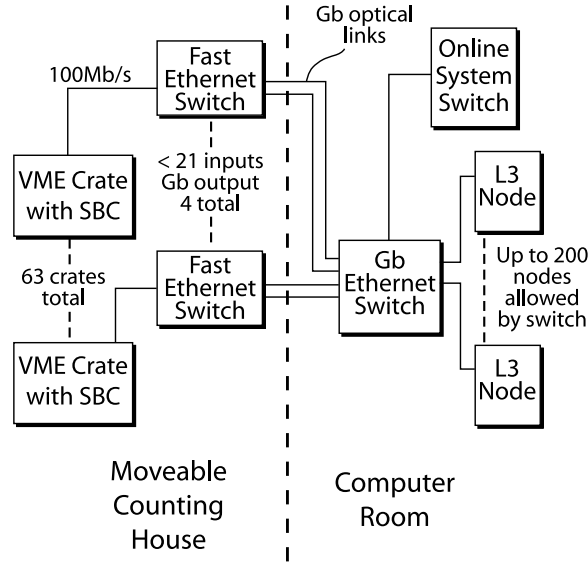


Fig. 71. The physical network configuration of the L3DAQ system.

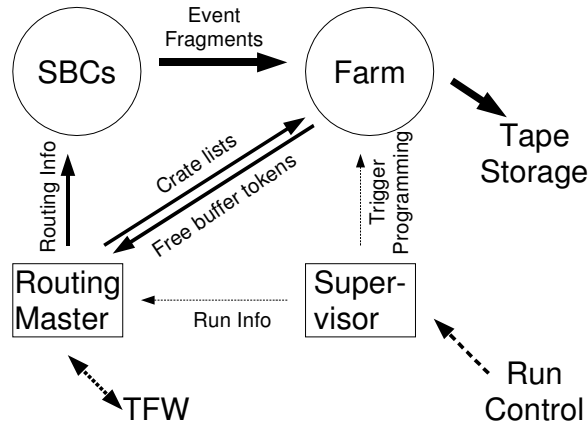


Fig. 72. Schematic illustration of the information and dataflow through the L3DAQ system.

The SBCs are single-board computers with dual 100 Mbyte/s Ethernet interfaces and a VME-to-PCI interface. An expansion slot is occupied by a digital-I/O (DIO) module, used to coordinate the readout of VME modules over the VME user (J3) backplane. A custom kernel driver on the SBC handles interrupt requests from the DIO module that are triggered by readout requests from the crate-specific electronics. On each readout request, the kernel module performs the VME data transfers and stores the event fragment in one of several buffers in kernel-memory.

A user-level process on the SBC receives route information from the RM in the form of route tags that contain a unique event identifier (L3 transfer number) and the indices of the farm nodes to which that event should be sent. If the route tag's L3 transfer number matches that of the transfer number embedded within the head event fragment in the kernel buffers, the event fragment is

sent to the specified farm nodes.

The EVB process on each farm node collates the event fragments received from SBCs into complete events, keyed by the L3 transfer number. For each event, the EVB receives an expected-crate list from the RM to determine when an event is complete. Complete events are placed in shared memory event buffers for processing by the L3 filter processes. The EVB routinely informs the RM of the number of available event buffers that it has. However, the EVB typically reports a smaller number of free buffers (currently a maximum of three) than it has available, ensuring that the total number of events, and thus the total volume of data, flowing through the main ethernet switch at any moment is smaller than the size of the output buffer memory of the switch.

The RM program executes on an SBC in a special VME crate containing a hardware interface to the TFW. The TFW provides trigger information and the transfer number for each event and allows the RM to asynchronously disable the firing of triggers. For each event, the RM program chooses a node for processing based on the run configuration, the trigger information, and the number of available buffers in the set of nodes configured to process the type of event. A node is chosen in a round-robin fashion from among the set of nodes with the most free buffers. If the number of available free buffers is too few, the RM instructs the TFW to disable triggers so that the farm nodes have time to catch up.

10.2 Online host

The online host system receives event data from the L3 farm nodes at a combined rate of approximately 12.5 Mbyte/s (50 Hz L3 accept rate of 250 kB events) and distributes that data to logging and monitoring tasks. The host system is capable of supporting multiple simultaneous data runs. Events that pass the L3 filters are tagged with a data stream identification that is a function of the satisfied hardware and software trigger components. Different streams are recorded independently; events are assigned to only one stream, excepting some events which are additionally assigned to special monitoring streams.

The final repository for the raw event data is tape, maintained in a robotic tape system located about 3 km from the detector. Data must be transmitted to each tape drive at approximately 10 Mbyte/s to keep the drive operating in streaming mode, since the remote tape drive nodes have no intermediate disk buffer. The online system is capable of simultaneous output to multiple tape streams and of buffering in case of tape robot unavailability. In addition to logging data, the online host system must supply between ten and twenty




Fig. 73. Schematic of the physical architecture of the online host system.

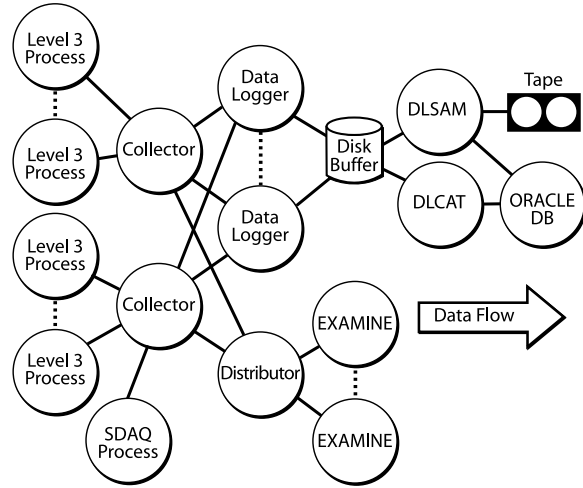


Fig. 74. Schematic of the software architecture of the online host system.

data monitoring clients at anywhere from 1% to 100% of the full data rate.

Figures 73 and 74 illustrate the physical and software architecture of the online host system. Event data arrive from the L3 trigger nodes at collector processes. The collector directs each event to the data logger appropriate for the stream identifier determined for the event. The collector also sends, on a best-effort basis (there is no flow control backpressure to the L3 nodes), a copy of each event to a distributor process, which is an event queueing system that provides event data in near real-time to analysis and monitoring clients (EXAMINE programs).

A diagnostic secondary data path, SDAQ, makes possible the processor-based readout of digitizing electronics. Information from SDAQ is also routed via the collector processes, allowing all downstream DAQ components to be shared. The SDAQ mechanism bypasses the L3DAQ and L3 trigger and is extremely valuable in the commissioning and calibration of detector components. The

detector-specific components of an SDAQ application have access to a library of SDAQ functions that handle queuing of data messages between components, interrupt management with callbacks, run synchronization, and priority-based scheduling. Two of the subdetectors use the SDAQ system: *i)* the SMT to monitor the performance of individual silicon detector channels during a run and *ii)* the CFT to calibrate the response of the fiber tracker channels.

The data logger writes data to files, grouped according to stream classification tags. Each data logger is responsible for a set of streams. The data logger also generates metadata information in file format for storage in a database to enhance the offline access of data. The DLSAM processes are the interfaces to the SAM/ENSTORE mass storage facility (Section 12.6 and the associated database descriptions). DLSAM monitors the local data buffer disks and requests that files be stored in the mass storage system (ENSTORE). This request is made through the database interface (SAM) which negotiates the request for file storage with the ENSTORE system.

All of the online computing systems are connected to a single high-capacity network switch. The architecture provides for parallel operation of multiple instances of the bandwidth-critical components.

The high-level online software applications are predominately constructed using the Python scripting language. Network communication between the components is implemented with the ITC package, a multi-platform, multi-threaded client/server messaging system based on the ACE [113] network and utility library.

11 Controls and Monitoring

The DØ experiment has extended EPICS (Experimental Physics and Industrial Control System) [114] to meet the control and monitoring requirements of a large high-energy physics detector. EPICS, an integrated set of software building blocks for implementing a distributed control system, has been adapted to satisfy the slow controls needs of the detector by *i)* extending the support for new device types and an additional field bus, *ii)* the addition of a global event reporting system that augments the existing EPICS alarm support, and *iii)* the addition of a centralized database with supporting tools for defining the configuration of the control system.

EPICS uses a distributed client-server architecture consisting of host-level nodes (clients) that run application programs and input/output controller (IOC) nodes (servers) that interface directly with the detector hardware. The two classes of nodes are connected by a local area network. Clients access

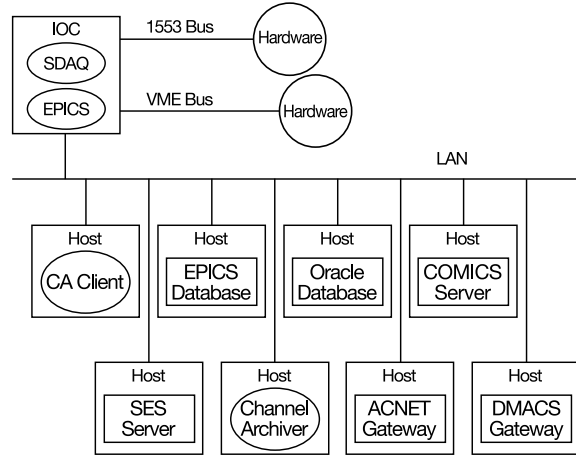


Fig. 75. Organization of the control system components. The Hosts are computers running Linux or Tru64 Unix and the IOCs are embedded computers in VME crates. Only one IOC is shown, and only a few of the online hosts.

process variable (PV) objects on the servers using the EPICS channel access protocol.

One of the unique properties of the DØ detector interface is the use of the MIL-STD-1553B serial bus for control and monitoring operations of the electronics components located in the collision hall. Since the detector is inaccessible for extended periods of time, a robust, high-reliability communication field bus is essential. EPICS was extended by providing a queuing driver for MIL-STD-1553B controllers and a set of device support routines that provide the adaptive interface between the driver and the standard EPICS PV support records. With these elements in place, all of the features of EPICS are available for use with our remote devices.

Figure 75 shows the architecture and components of the DØ controls and monitoring system, a distributed system built using EPICS.

11.1 Global Event Reporting

To process significant events from all experiment sources, a separate facility, the significant event system (SES), collects and distributes all changes of state. The SES has a central server that collects event messages from sender clients and filters them, via a Boolean expression, to route to receiving clients. Sender clients, including the IOCs, connect to the server and all state changes on those clients, including alarm transitions, are sent to the server. The architecture of the SES and the flow of messages within the system are illustrated in Figure 76.

The alarm class of SES messages receives special handling in the server. The SES server maintains the current alarm state of the entire experiment so that

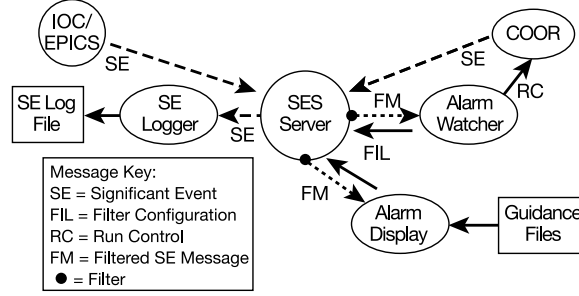


Fig. 76. The major components of the SES are the server, logger, watcher, and alarm display. COOR is a coordinating process that provides run control to the experiment. Significant events originate either with the detector hardware or the run control system.

receiving clients can obtain the current state when they first connect to the server. In addition to specialized receiving clients that may connect to the server, there are three standard clients: the SES logger, the SES alarm display, and the SES alarm watcher. The logger has a pass-all filter so that it receives all SES messages sent to the server and writes the messages received to a disk file. The current state of the detector stored in the server is relayed to users through the alarm display. For alarms that compromise data quality, the alarm watcher automatically pauses the current run. In addition to its monitoring and logging functions, the SES system provides the means for distributing synchronizing messages to other components of the online software system. Tools have been developed for mining data from the SES log files. Hardware experts review the log files to understand which hardware devices are unstable and collaborators performing data analysis can ensure the event they are examining is real and not caused by a fault in the detector.

11.2 Centralized device database

The EPICS databases that configure the individual IOCs are flat ASCII files of record definitions, the database equivalent of a PV, that are read by the IOCs during startup. The EPICS system additionally provides a higher-level construct, called a template, which is a parameterized collection of record definitions. Generator files, which reference the templates, supply the parameter values to produce instances of these templated devices. While these collections of files are adequate for EPICS initialization, they are not easily accessible to host-level processes which require the same information.

To address this problem, we centralized the relevant information in a relational database (Oracle [115]) and provided a family of scripts to manage the transformation between the relational database and the EPICS ASCII-format files. At the time this document was prepared, the database contained approximately 6200 templated devices, corresponding to about 137000 process

variables, and that number is constantly expanding

In addition to the database management scripts, a WWW browser interface to the relational database is available for the initial definition, modification, and viewing of the relational database entries. With control system device specifications centralized in the relational database, they are easily accessible to other host-level processes.

11.3 Detector configuration management

One of the most complex tasks performed by the control system is the configuration of the detector for specific run conditions. The set of distinct configurations, both for normal data-taking and for calibration runs, is very large; the usual technique of uploading a specific detector configuration and saving it as a file for subsequent downloading is impractical.

For ease of configuration management, the detector is represented as a tree with nodes at successively deeper levels corresponding to smaller, more specialized organizational units of the detector. The terminal nodes of the tree, called action nodes, each manage the configuration of a specific, high-level device. They are instances of the high-level devices discussed in the preceding database section. The intermediate nodes of the tree primarily serve to organize the traversal order of the subordinate nodes since the detector is, in general, sensitive to the order in which devices are initialized.

A single server program, COMICS, manages configuration of the EPICS-accessible part of the detector. The tree nodes, both intermediate and action, are all specialized instances of a base node class that defines the majority of the methods that characterize node behavior. The detector tree structure is defined by a set of configuration files that are Python program segments which instantiate instances of these nodes.

11.4 Operator Interfaces and Applications

An application framework, in the Python scripting language, assists in developing operator interfaces and provides a consistent “look and feel” for all visual displays. This framework includes a collection of specialized, graphical objects for constructing updating displays of PV values using a Python interface to the EPICS channel access API (application program interface). The experiment uses more than forty instances of these monitoring displays in the control room to manage the detector components.

11.5 Archiving EPICS data

While extensively using EPICS records for control and monitoring tasks, almost every detector group needs to maintain structured access to archived PV values. There are two major archiving tools employed by DØ: (1) the channel archiver [116], for needs that require rather fast sampling rates (1 Hz or faster) but do not require frequent access to historical data; and (2) the EPICS/Oracle Archiver, for long-term studies that require slow sampling rates (once per minute or less frequently), easy access to data at any moment, and minimal maintenance.

Many channel archivers are running concurrently, monitoring several thousand PVs. About once a week, collected archives are sent to the central Fermilab robotic tape storage via the SAM data management system (Section 12.6). The channel archiver toolset has many user-friendly interfaces, including web-based tools, that enable retrieval from an archive of data in different formats and generation of time plots with various options.

11.6 ACNET gateway

In the operation of the detector, it is vital to have a fast and reliable messaging connection between DØ and accelerator operations to exchange control and monitoring information. The DØ control system supplies cryogenic and magnet data and FPD pot positions and counter rates. The accelerator control system (ACNET), in turn, sends information about critical accelerator devices. A gateway between the DØ and ACNET control systems, based on the XMLRPC protocol, provides this interconnection.

12 Computing and Software

A large amount of software has been developed for data acquisition, for monitoring and controlling hardware, for Monte Carlo event simulation, and for data and Monte Carlo event reconstruction. Early in the development of software for Run II, we made the decision that all new software would be written using the C++ programming language. Legacy Run I Fortran software and programs from other sources (e.g. Monte Carlo event simulation programs) are wrapped in C++ code. In this section, we give an overview of the computing and software used during Run II.

12.1 *Event data model*

The DØ event data model (EDM) is a library of C++ classes and templates whose purpose is to support the implementation of reconstruction and analysis software.

The central feature of the EDM is the event, a class that represents the results of a single beam crossing. The event acts as a container to manage all the data associated with a single crossing: the raw output of the detector, the results of L3 filter processing, and the results of many different reconstruction tasks. Each of the items in this collection contains both the data describing the crossing in question as well as meta-data that describe the configuration of the program that constructed these results. This allows us to run multiple instances of single algorithms with different configuration information (for example, several cone-based jet algorithms with different cone radii), and to distinguish between the output of these different algorithms. The EDM also provides a mechanism for access to the collected reconstruction results, relieving the users from the burden of understanding the somewhat complex internal organization required for the management of the event data and the corresponding meta-data.

12.2 *Data persistency*

The conversion of the C++ objects used in the reconstruction program to a persistent format is handled by the DØ object model (DØOM) [117]. This has several parts (Figure 77). First, DØOM maintains a dictionary describing the layout of the C++ classes that are to be used persistently, which is generated by running a preprocessor over the C++ headers defining the classes. This preprocessor is based on a modified version of the CINT C/C++ interpreter (which is also used in the ROOT system [118]).

The actual translation between C++ objects and the persistent format is handled by one of several I/O packages. User code has no dependence on the I/O packages, so that new formats can be added without changing any reconstruction code. The I/O packages are typically built on top of external software packages to do the actual I/O.

Finally, the external interface to the package is provided by a set of stream classes. A set of DØ framework packages that use these classes to read and write events within the framework is also provided; in most cases, a framework user need only to include these packages in order to read and write events.

DØOM includes numerous features to assist with schema evolution and ver-

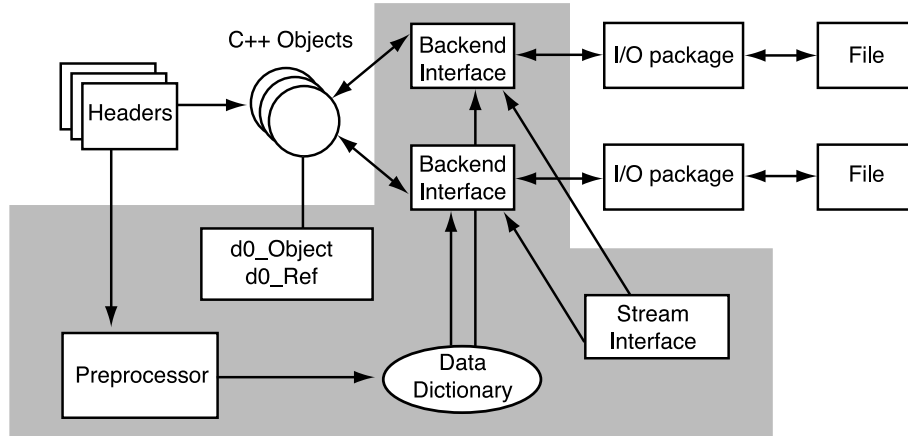


Fig. 77. The structure of DØOOM. The shaded area shows the components which are part of DØOOM proper.

sioning of data. The dictionary information is maintained along with the saved data, so the layout of the saved data is always known. During reading, class members are matched between the C++ and persistent forms based on name. This allows DØOOM to handle the common cases of adding and deleting data members with no explicit action on the programmer's part. For more complicated cases, it is possible to provide conversion code that is automatically run when needed. The DØOOM dictionary information may also be queried at run time. The system can also make use of the persistent dictionary information to build objects at run time for which no dictionary information was compiled into the program. This is useful for programs to dump or browse arbitrary DØOOM data files.

One I/O package is based on the DSPACK library, which originated with the NA49 experiment [119]. DSPACK handles the conversion from C-like structures to a serial data format; the DØOOM I/O package converts from the C++ objects to the DSPACK structures. At DØ, DSPACK data are usually encapsulated inside another, lightweight, format called EVPACK, which provides data compression and random access within files with keyed lookup. EVPACK-encapsulated DSPACK records may also be sent over the network; this is used to distribute data within the online host system. All event data are stored in (EVPACK-encapsulated) DSPACK format. In addition, this format is used for several static, structured data files used by the reconstruction program, such as the description of the detector geometry.




Fig. 78. The hierarchy of database, middle tier server, and client applications. Each server can cache data, and the servers can be chained to provide higher-speed access for clients located at remote processing sites.

12.3 Calibration databases

Run-dependent information needed for reconstruction, such as magnetic field polarities and calibration constants, is stored in a database. To decouple the reconstruction program from any particular proprietary database implementation, DØ makes use of a three-tier architecture [120], where the database and the client code make up the top and bottom tiers (Figure 78). Between is a middle tier server which accepts requests from the clients, makes the query to the database, and returns the data to the client. Data can be cached by the server, thus speeding up the response to the client and reducing the load on the database.

This caching is especially effective for the large-scale reconstruction of raw data. If runs that require the same calibration constants are processed together, only one query to the database is made for each calibration set; all other requests for this set are fulfilled directly from the cache.

Servers can be chained together so that processing sites outside Fermilab may run their own local servers which request their data from the servers at the laboratory. This allows a remote processing site to have its own cache to speed up queries without requiring the site to access the database directly.

The particular software used for storing the calibration data is the Oracle relational database. The middle tier servers are mostly implemented in the

Python programming language with some parts written in C++ to improve performance. CORBA distributed object technology [121] is used for the communication between the client and the servers. The core part of the server is written to be independent of the particular application it is serving data for; the customized code which queries specific database tables and returns the results to the client is automatically generated from the database schema for each particular application. Each subsystem of the detector (SMT, CFT, preshower detectors, calorimeter, and muon detectors) has its own database schema and a separate server.

The client end of the system is an I/O package for DØOM. The necessary C++ class definitions are also automatically generated from the database schema. Accessing the data is then a matter of navigating a C++ object hierarchy: the client application need have no knowledge of CORBA, Oracle, or the method by which it obtained the data. This decoupling makes it relatively easy to replace the specific implementation of the middle tier, if this becomes either necessary or desirable, without requiring the whole of the client implementation to change.

12.4 Event simulation

The generation of Monte Carlo (MC) events involves multiple stages and many executables. To integrate all processes, all programs use the EDM to carry data in memory and DØOM to store persistent event data. All code is organized in independent packages running in a standard DØ framework and is written in C++ or embedded in C++ driving routines.

The first step in MC event generation is the simulation of a physical process, a $p\bar{p}$ collision producing a particular final state. Nearly all existing event generator programs are written in Fortran, but the StdHep code from the FNAL Computing Division can be used to store the output in a standard common block format. This allowed us to write a C++ wrapper that converts the StdHep Fortran format to C++ classes satisfying the EDM requirements.

To trace the particles through the DØ detector, determine where their paths intersect active areas, and simulate their energy deposition and secondary interactions, we use the CERN program GEANT v3.21, which is also written in Fortran. A C++ wrapper is used to read files produced by the event generators and to write the output of GEANT in DØOM format. This executable is called DØGSTAR. All subsequent steps in the event simulation are handled by programs written almost entirely in C++.

The DØSIM program modifies the generated Monte Carlo data to account for various detector-related effects. After the particles from the simulated reaction

have been traced through the detector, the generated energy depositions must be converted to the form that the real data takes when processed through the DØ electronics. Detector inefficiencies and noise (from the detector and electronic readout) must be taken into account, and more than one interaction may occur during a beam crossing. In addition, some portions of the detector (such as the calorimeter) remain sensitive to interactions over a period of time that includes more than one beam crossing. Simulation of the trigger electronics and the effects of the trigger on data selection is performed by a separate program, DØTRIGSIM. DØTRIGSIM contains simulation code only for the L1 trigger. The L2 and L3 triggers are based on filtering code, and exactly the same software runs in DØTRIGSIM. The output of DØSIM and DØTRIGSIM is in the same format as the data recorded by the DØ data acquisition system, but contains additional MC information to make it possible to correlate detector data with the original generator output.

12.5 *Reconstruction*

The DØ offline reconstruction program (RECO) is responsible for reconstructing objects used for physics analysis. It is a CPU-intensive program that processes collider events recorded during online data taking and simulated MC events. The executable is run on the offline production farms and the results are placed into the central data storage system (Section 12.6) for further analysis. Information and results for each event are organized using the EDM. The EDM manages information within the event in blocks called chunks. The raw data chunk (RDC), created either by an L3 processor node or the MC, contains the raw detector signals and is the primary input to RECO. The output from RECO consists of many additional chunks associated with each type of reconstructed object. RECO produces two output formats or data tiers. The data summary tier (DST) contains all information necessary to perform any physics analysis, including limited re-reconstruction of high-level physics objects. The thumbnail (TMB) is a physics summary format less than one-tenth the size of the DST format. The TMB can be used directly to perform many analyses, and it allows the rapid development of event selection criteria to be applied to the DST sample.

RECO reconstructs events in several hierarchical steps. The first involves detector-specific processing. Detector unpackers process individual detector data blocks within the RDC, decoding the raw information, associating electronics channels with physical detector elements, and applying detector-specific calibration constants. For many of the detectors, this information is then used to reconstruct cluster (for example, from the calorimeter and preshower detectors) or hit (from the tracking detectors) objects. These objects use geometry constants to associate detector elements (energies and positions) with phys-

ical positions in space. The second step in RECO focuses on the output of the tracking detectors, reconstructing global tracks from the hits in the SMT and CFT. This process, involving several different tracking algorithms, is the most CPU-intensive activity of RECO. The results are stored in corresponding track chunks, which are used as input to the third step of RECO, vertexing. First, primary vertex candidates are found. These vertices indicate the locations of $p\bar{p}$ interactions and are used in the calculation of various kinematical quantities (e.g. E_T). Next, displaced secondary vertex candidates are identified. Such vertices are associated with the decays of long-lived particles. The results of the above algorithms are stored in vertex chunks, and are then available for the final step of RECO — particle identification. Using a wide variety of algorithms, information from each of the preceding reconstruction steps is combined and standard physics object candidates are created. RECO first finds electron, photon, muon, neutrino (\not{E}_T), and jet candidates, after which it identifies candidates for heavy-quark and tau decays are identified next.

12.6 Data handling and storage

The sequential access via meta-data [122] (SAM) data handling system gives users access to all the data created by the DØ experiment (both detector data and simulation data), in a very flexible and transparent manner: the user does not need to know where the files are physically stored, nor worry about exactly how they are delivered to her/his process. SAM oversees the functions of cataloging data (files and events, and associated metadata regarding production conditions), transferring data in and out of mass storage systems, transferring data among different computer systems (whether connected via local or wide area network), allocating and monitoring computing resources (batch slots, tape mounts, network bandwidth, disk cache space), and maintaining file delivery status at the user process level. The bookkeeping functions of the SAM system are provided by an Oracle [115] database, which is accessed via a client-server model utilizing CORBA technology. Files are stored in SAM using interfaces that require appropriate metadata for each file. The files are organized, according to the metadata provided, by data tier, and by production information (program version which produces the data, etc.). The SAM system also provides file storage, file delivery, and file caching protocols that permit the experiment to control and allocate the computing resources. Tape resources can be guaranteed to high priority activities (data acquisition and farm reconstruction), high usage files can be required to remain in the disk cache, and different priorities and allocations for resource usage can be granted to groups of users.

Acknowledgements

We would like to thank Ken Ford and Scott Baxter for their work on the illustrations and to acknowledge in-kind contributions by the Altera and Xilinx Corporations. We thank the staffs at Fermilab and collaborating institutions, and acknowledge support from the Department of Energy and National Science Foundation (USA), Commissariat à l'Energie Atomique and CNRS/Institut National de Physique Nucléaire et de Physique des Particules (France), Ministry of Education and Science, Agency for Atomic Energy and RF President Grants Program (Russia), CAPES, CNPq, FAPERJ, FAPESP and FUNDUNESP (Brazil), Departments of Atomic Energy and Science and Technology (India), Colciencias (Colombia), CONACyT (Mexico), KRF (Korea), CONICET and UBACyT (Argentina), The Foundation for Fundamental Research on Matter (The Netherlands), PPARC (United Kingdom), Ministry of Education (Czech Republic), Natural Sciences and Engineering Research Council and WestGrid Project (Canada), BMBF and DFG (Germany), A.P. Sloan Foundation, Civilian Research and Development Foundation, Research Corporation, Texas Advanced Research Program, and the Alexander von Humboldt Foundation.

References

- [1] S. Abachi, et al., Phys. Rev. Lett. 74 (1995) 2632.
- [2] S. Abachi, et al., Phys. Rev. Lett. 79 (1997) 1197.
- [3] B. Abbott, et al., Phys. Rev. Lett. 80 (1998) 2063.
- [4] B. Abbott, et al., Phys. Rev. D 58 (1999) 052001.
- [5] B. Abbott, et al., Phys. Rev. D 60 (1998) 052001.
- [6] V. Abazov, et al., Nature 429 (2004) 638.
- [7] S. Abachi, et al., Phys. Rev. Lett. 77 (1996) 3309.
- [8] B. Abbott, et al., Phys. Rev. D 58 (1998) 12002.
- [9] B. Abbott, et al., Phys. Rev. Lett. 80 (1998) 3008.
- [10] B. Abbott, et al., Phys. Rev. D 58 (1998) 092003.
- [11] B. Abbott, et al., Phys. Rev. D 62 (2000) 092006.
- [12] B. Abbott, et al., Phys. Rev. Lett. 84 (2000) 222.
- [13] V. M. Abazov, et al., Phys. Rev. D 66 (2002) 012001.

- [14] S. Abachi, et al., Phys. Rev. Lett. 78 (1997) 3634.
- [15] S. Abachi, et al., Phys. Rev. Lett. 78 (1997) 3640.
- [16] S. Abachi, et al., Phys. Rev. D 56 (1997) 6742.
- [17] B. Abbott, et al., Phys. Rev. Lett. 79 (1997) 1441.
- [18] B. Abbott, et al., Phys. Rev. D 57 (1998) 3817.
- [19] B. Abbott, et al., Phys. Rev. D 58R (1998) 051101.
- [20] B. Abbott, et al., Phys. Rev. D 58R (1998) 31102.
- [21] B. Abbott, et al., Phys. Rev. D 60 (1999) 072002.
- [22] B. Abbott, et al., Phys. Rev. D 62 (2000) 052005.
- [23] B. Abbott, et al., Phys. Rev. Lett. 82 (1999) 2451.
- [24] B. Abbott, et al., Phys. Lett. B 525 (2002) 211.
- [25] B. Abbott, et al., Phys. Rev. Lett. 82 (1999) 2451.
- [26] B. Abbott, et al., Phys. Rev. D 64 (2001) 032003.
- [27] B. Abbott, et al., Phys. Rev. Lett. 79 (1997) 4321.
- [28] B. Abbott, et al., Phys. Rev. Lett. 80 (1998) 2051.
- [29] B. Abbott, et al., Phys. Rev. Lett. 81 (1998) 38.
- [30] B. Abbott, et al., Phys. Rev. Lett. 83 (1999) 2896.
- [31] B. Abbott, et al., Phys. Rev. Lett. 84 (2000) 2088.
- [32] V. M. Abazov, et al., Phys. Rev. D 64 (2001) 092004.
- [33] V. M. Abazov, et al., Phys. Rev. Lett. 8 (2002) 191801.
- [34] S. Abachi, et al., Phys. Rev. D 57 (1998) 589.
- [35] B. Abbott, et al., Phys. Rev. Lett. 80 (1998) 442.
- [36] B. Abbott, et al., Phys. Rev. Lett. 80 (1998) 1591.
- [37] B. Abbott, et al., Phys. Rev. Lett. 82 (1999) 29.
- [38] B. Abbott, et al., Phys. Rev. D 60R (1999) 031101.
- [39] B. Abbott, et al., Phys. Rev. Lett. 83 (1999) 4937.
- [40] V. M. Abazov, et al., Phys. Rev. Lett. 88 (2002) 171802.
- [41] V. M. Abazov, et al., Phys. Rev. D 66 (2002) 112001.
- [42] See www-d0.fnal.gov/www_buffer/pub/publications.html for a complete list of DØ publications.

- [43] S. Abachi, et al., Nucl. Instrum. Methods Phys. Res. A 338 (1994) 185.
- [44] Micron Semiconductor Ltd., 1 Royal Buildings, Marlborough Road, Lancing, Sussex, BN15 8UN, England.
URL www.micronsemiconductor.co.uk
- [45] Canberra Eurisys, 4, avenue des Frênes, 78067 St. Quentin Yvelines CEDEX, France.
URL www.eurisysmesures.com
- [46] ELMA: Research and Production Organization for Electronic Materials (NPO ELMA), 103460 Moscow, Russia.
- [47] M. Moll, Radiation Damage in Silicon Particle Detectors — Microscopic Defects and Macroscopic Properties, Ph.D. thesis, University of Hamburg, dESY-THESIS-1999-040, ISSN 1435-8085 (1999).
- [48] DuPont High Performance Materials, P.O. Box 89, Circleville, OH 43113.
URL www.dupont.com/kapton
- [49] DuPont High Performance Materials, P.O. Box 89, Circleville, OH 43113.
URL www.dupont.com/tedlar
- [50] Emerson & Cuming, 46 Manning Road, Billerica, MA.
URL www.emersoncuming.com
- [51] Hirose Electric U.S.A., 2688 Westhills Ct., Simi Valley, CA 93065.
URL www.hiroseuse.com
- [52] R. Yarema, et al., Fermilab-TM-1892 (1994, revised 1996).
- [53] United Technologies Microelectronics Center, Inc. is now part of Aeroflex.
- [54] R. Ruchti, Ann. Rev. Nucl. Part. Sci. 46 (1996) 281.
- [55] Kuraray International Corporation, 200 Park Ave., New York, NY 10166.
URL www.kuraray-am.com
- [56] Degussa AG, Dusseldorf, Germany.
URL www.rohacell.com
- [57] GE Polymerland, Inc.
URL www.gepolymerland.com
- [58] Crompton Corporation.
URL www.cromptoncorp.com
- [59] M. D. Petroff, M. G. Stapelbroek, IEEE Trans. Nucl. Sci. 36 (1989) 158.
- [60] M. D. Petroff, M. Atac, IEEE Trans. Nucl. Sci. 36 (1989) 163.
- [61] A. Bross, other, Nucl. Instrum. Methods Phys. Res. A 477 (2002) 172.
- [62] R. A. Rucinski, Short Profile Rectangular Helium Cryostat Made From Nickel-Iron Alloy, in: S. Breon, et al. (Eds.), Advances in Cryogenic Engineering, Vol. 46B, American Institute of Physics, Melville, New York, 2002, p. 1468.

- [63] CINCH Connector Division, A Division of Labinal Components & Systems, Inc., 1501 Morse Ave., Elk Grove Village, IL 60007.
URL www.cinchuk.com/cinapse.htm
- [64] Texas Instruments, Inc.
URL focus.ti.com/lit/ds/symlink/sn65lvds95.pdf
- [65] Microchip Technology, Inc., 2355 West Chandler Blvd., Chandler, AZ 85224-6199.
URL www.microchip.com
- [66] B. Baumbaugh, et al., IEEE Trans. Nucl. Sci. 45 (1998) 343.
- [67] Nichia Corp., 491 Oka, Kaminaka-Cho, Anan-Shi, Tokushima 774-8601, Japan.
Tel. +81 (884) 22-231.
- [68] A. M. Patwa, The Forward Preshower System and a Study of the J/ψ Trigger with the DØ Detector, Ph.D. thesis, SUNY at Stony Brook (2002).
URL www-d0.fnal.gov/results/publications_talks/thesis/patwa/thesis.ps
- [69] J. Brzenziak, et al., Fermilab-TM-1886 (1994).
- [70] Toshiba Corp, 1-1, Shibaura 1-chome, Minato-ku, Tokyo 105-8001, Japan.
URL www.toshiba.co.jp
- [71] J. Simkin, C. Trowbridge, Three-Dimensional Computer Program (TOSCA) for Non-linear Electromagnetic Fields, RL-79-097 (1979).
- [72] P. Baringer, et al., Nucl. Instrum. Methods Phys. Res. A 469 (2001) 295.
- [73] Abatron Inc., 33 Center Dr., Gilberts, IL 60136.
- [74] Altera Corporation, 101 Innovation Drive, San Jose, CA 95134, document part number DS-F10K-4.2.
URL www.altera.com
- [75] A. Besson, Étude des événements di-leptons + 4 jets dans le Run II del l'expérience D0 à Fermilab, Ph.D. thesis, Université de Grenoble 1, iSN 02-50 (2002).
URL www-d0.fnal.gov/results/publications_talks/thesis/besson/these.ps
- [76] Xilinx, Inc., 2100 Logic Drive, San Jose, CA 95124-3400.
URL www.xilinx.com
- [77] National Instruments Corporation, 11500 N. Mopac Expwy., Austin, TX 78759-3504.
URL www.labview.com
- [78] Saint-Gobain Crystals & Detectors.
URL www.bicron.com

- [79] Hamamatsu Corporation.
URL www.hamamatsu.com
- [80] B. Acharya, et al., Nucl. Instrum. Methods in Phys. Res. A 401 (1997) 45.
- [81] Glasteel Tennessee, Inc., 175 Commerce Rd., Collierville, TN 38017.
URL www.glasteel.com
- [82] Dupont, Inc.
URL tyvek.com
- [83] THORN EMI, Gencom Inc, 23 Madison Rd., Fairfield, NJ 07006.
- [84] MELZ Company, Elektrozavodskaya str., 23, Moscow Russia, 105023.
- [85] V. Evdokimov, Light Collection from Scintillation Counters Using WLS Fibers and Bars, in: Proceedings of SCIFI'97, American Institute of Physics, Melville, New York, 1997, p. 300.
- [86] R. Veenhof, Simulation of gaseous detectors.
URL consult.cern.ch/writeup/garfield/
- [87] S. Belikov, et al., Instruments and Experimental Technique 36 (1993) 390.
- [88] J. Butler, et al., Fermilab-fn-629 (1995).
- [89] N. Mokhov, Fermilab-fn-628 (1995).
- [90] N. Mokhov, Fermilab-conf-03/053 (2003).
URL www-ap.fnl.gov/MARS
- [91] P. Hanlet, et al., Nucl. Instr. and Methods in Phys. Res. A 521 (2004) 343.
- [92] A. Brandt, et al., Fermilab-Pub-97/377 (1997).
- [93] U. Amaldi, et al., Phys. Lett. B43 (1973) 231.
- [94] Phillips Scientific.
URL www.phillipsscscientific.com
- [95] C. Miao, Nucl. Phys. Proc. Suppl. 78 (1999) 342.
- [96] A. Pla-Dalmau, A. Bross, C. Hurlbut, fermilab-Conf-94/096 (1994).
- [97] H. Corporation, Photomultiplier tubes, basics and applications, 2nd edition.
- [98] Times Microwave Systems LMR 240 cables are used to bring the photomultiplier signals out of the detector and Times Microwave Systems LMR 400 cables are used to bring the signals the rest of the way to the Moving Counting House.
URL www.timesmicrowave.com
- [99] G. Gao, R. Partridge, IEEE Trans. Nucl. Sci. 38 (1991) 286.
- [100] M. Abolins, et al., Nucl. Instr. and Methods A289 (1990) 542–560.

- [101] J. Anderson, et al., IEEE Trans. Nucl. Sci. 51 (2004) 345–350.
- [102] S. Rapisarda, N. Wilcer, R. Yarema, IEEE Trans. Nucl. Sci. 50 (2003) 878.
- [103] C.-S. Yen, G-link: A chipset for gigabit-rate data communication (1992).
- [104] Cypress Semiconductor Corporation.
URL www.cypress.com
- [105] VME64 Extensions for Physics and Other Applications (VME64xP), VITA 23-1998, Draft 1 (February 1999).
- [106] The Mbus was originally designed for use by the CDF experiment. DØ uses a similar physical backplane, but implements a re-engineered communications protocol to support effective arbitration in an environment with numerous master and target devices.
- [107] J. M. C. Murphy, M. Campbell, in: G. J. Blunar, R. L. Sumner (Eds.), Proceedings of the 4th International Conference on Electronics for Future Colliders, LeCroy, Chestnut Ridge, NY, 1994, p. 101.
- [108] R. Hirosky, in: S. Schaller (Ed.), Proceedings of the 11th IEEE Real Time Conference 99, IEEE, 1999, p. 512.
- [109] PICMG 2.0 R2.1, CompactPCI specification, pCI Industrial Manufacturers Group (1997).
URL www.picmg.org
- [110] Tundra Semiconductor Corporation, 603 March Road, Ottawa, Ontario, Canada K2K 2M5, Document 8091142.MD300.01.
URL www.tundra.com
- [111] M. Fortner, et al., IEEE Trans. Nucl. Sci. 49 (2002) 1589–1592.
- [112] Cisco Systems, Inc.
URL www.cisco.com
- [113] Douglas C. Schmidt and others.
URL www.cs.wustl.edu/~schmidt/ACE.html
- [114] L. Dalesio, et al., The experimental physics and industrial control system architecture: Past, present, and future, in: Proceedings of ICALEPCS93, Berlin, Germany, 1993, pp. 179–184.
- [115] Oracle Corporation, 500 Oracle Parkway, Redwood Shores, CA 94065.
URL www.oracle.com
- [116] K. Kazemir, L. Dalesio, Data Archiving in EPICS, in: Proceedings of ICALEPCS99, Trieste, Italy, 1999, p. ??
- [117] S. Snyder, The DØ Object Model, presented at Computing in High Energy Physics '98, Chicago (1998).
URL www.hep.net/chep98/PDF/130.pdf

- [118] R. Brun, et al., ROOT: An Object-Oriented Data Analysis Framework.
URL root.cern.ch/
- [119] R. Zyburt, P. Buncic, DSPACK: Object Manager for High Energy Physics, in:
Proceedings of Computing in High Energy Physics '95, Rio de Janeiro, 1995,
pp. 345–348.
- [120] H. Greenlee, et al., Serving Database Information Using a Flexible Server in
a Three-Tier Architecture, presented at Computing in High Energy Physics
2003 (2003).
URL www.arxiv.org/abs/cs.DC/0307001
- [121] Object Management Group, The Common Object Request Broker:
Architecture and Specification (1998).
URL www.omg.org
- [122] A. Baranovski, et al., Fermilab-TM-2175 (2002).